



Research Institute for Advanced Computer Science
NASA Ames Research Center

RIACS ANNUAL REPORT

October 1, 1997 through September 30, 1998

Cooperative Agreement: NCC 2-1006

Submitted to:

Contracting Office Technical Monitor:
Anthony R. Gross
Code I, MS 200-3

Submitted by:

Research Institute for Advanced Computer Science (RIACS)

**An Institute of
Universities Space Research Association (USRA)**

RIACS Principal Investigator
Robert C. Moore, Director

Robert C. Moore, Director

TABLE OF CONTENTS

PARAGRAPH	TITLE	PAGE
I.	INTRODUCTION	2
II.	RESEARCH PROJECTS	4
	A. Autonomous Reasoning	4
	B. Human Center Computing	10
	C. High Performance Computing and Networks	20
III.	TECHNICAL REPORTS	34
IV.	PUBLICATIONS	41
V.	PAPERS SUBMITTED TO REFERRED JOURNALS	45
VI.	SEMINARS AND COLLOQUIA	47
VII.	OTHER ACTIVITIES	52
VIII.	RIACS STAFF	58

I INTRODUCTION

Robert C Moore, Director

The Research Institute for Advanced Computer Science (RIACS) was established by the Universities Space Research Association (USRA) at the NASA Ames Research Center (ARC) on June 6, 1983. RIACS is privately operated by USRA, a consortium of universities that serves as a bridge between NASA and the academic community. Under a five-year co-operative agreement with NASA, research at RIACS is focused on areas that are strategically enabling to the Ames Research Center's role as NASA's Center of Excellence for Information Technology. Research is carried out by a staff of full-time scientist, augmented by visitors, students, post doctoral candidates and visiting university faculty.

The primary mission of RIACS is charted to carry out research and development in computer science. This work is devoted in the main to tasks that are strategically enabling with respect to NASA's bold mission in space exploration and aeronautics. There are three foci for this work:

Automated Reasoning
Human-Centered Computing
High Performance Computing and Networking

RIACS has the additional goal of broadening the base of researcher in these areas of importance to the nation's space and aeronautics enterprises. Through its visiting scientist program, RIACS facilitates the participation of university-based researchers, including both faculty and students, in the research activities of NASA and RIACS. RIACS researchers work in close collaboration with NASA computer scientists on projects such as the Remote Agent Experiment on Deep Space One mission, and Super-Resolution Surface Modeling.

In 1998, RIACS had 16 research scientists, 14 visiting scientists, 3 postdoctoral scientists, 2 students and 3 administrators.

In collaboration with NASA, RIACS supported two Computation Fluid Dynamic short courses given by Professor Eli Turkel, Department of Mathematics, and Tel Aviv University. The courses were held to discuss need and opportunities in basic research in computer science in and for NASA applications.

Topic: Preconditioning Methods for Multidimensional Aerodynamics
Date: February 6, 1998

The purpose of the seminar was to discuss accuracy considerations and acceleration techniques for the numerical solution to the steady state compressible Navier-Stokes equations. This included discussions of (i) scalar, matrix and CUSP artificial viscosities for use with central difference approximations, (ii) preconditioning methods for low speed applications both with respect to convergence acceleration and accuracy, (iii) Jacobi preconditioning for damping high frequency errors with multi-grid methods, (iv) nonlocal preconditioners, (v) a presentation of some properties of the Navier-Stokes equations and their numerical approximation and possible future applications including applications to all speeds.

Topic: Numerical methods for Aeroacoustics and CEM
Date: February 10, 1998

The purpose of the seminar was to discuss the use of high order methods for the time dependent fluid, dynamic, compressible, equations, and also for the time dependent Maxwell equations. Including concentration on compact implicit methods of various kinds including extensions of the MacCormack scheme. Other discussions included differences between the total error and amplitude/phase errors, boundary conditions on PML (Perfectly Matched Layer) nonreflecting layers and extensions to aeroacoustics.

Furthermore, RIACS scientists participated and co-sponsored the American Association for Artificial Intelligence (AAAI-98) in Madison, Wisconsin.

RIACS technical reports are usually preprints of manuscripts that have been submitted to research journals or conference proceedings. A list of these reports for the period October 1, 1997 through September 30, 1998 is in the Reports and Abstracts section of this report.

II. RESEARCH PROJECTS

A. AUTONOMOUS REASONING

REMOTE AGENT EXPERIMENT (RAX)

Pandurang Nayak, Ph.D. (RIACS)

Research over the last year focused primarily in the areas of building model-based autonomous systems and the DS-1 Remote Agent Experiment. In the past RIACS co-developed the Livingstone model-based diagnosis and recovery system. Livingstone is a kernel of a reactive model-based autonomous system that performs significant deduction within the reactive control loop. In the past year, RIACS initiated research on extending Livingstone in a number of different ways. First, work was performed on extending Livingstone to be a ground operations station, that provides explanations for its reasoning, what-if analysis, and ambiguity management. The proposed solutions take into account the increased computational resources available on the ground and the existence of a human operator in the loop. Second, work was performed to extend Livingstone to be a complete on-board executive. This requires developing a tight coupling between Livingstone and Burton, a model-based planner developed earlier. Third, work was done to develop a model-based programming environment that would enable domain experts, rather than AI experts, to develop and deploy Livingstone models. Research in all three areas is expected to continue into the next fiscal year.

The primary focus of the work on the DS-1 Remote Agent Experiment (RAX) involved integration and test. Testing of an agent that synthesizes its own sequences and fault protection responses are particularly challenging since exhaustive testing and analysis is infeasible. We developed a scenario-based testing methodology that enabled us to adequately test the RAX within severe time and resource constraints. Research in the context of the RAX also resulted in a series of papers describing the over all Remote Agent architecture and a novel hybrid procedural/deductive executive (see below). Recent research, motivated by the computational constraints faced by RAX, has resulted in new algorithms that speed up Livingstone's diagnostic capability by a factor of 2-3.

THE EXTRACTION OF THREE DIMENSIONAL SUPER RESOLVED INFORMATION FROM MULTIPLE IMAGES OF THE SAME REGION.

Robin Morris, Ph.D. (RIACS)

This is an ongoing project. The primary goal is to build an inference system that will enable the information contained in a number of different views of the same scene to be combined into a single, high resolution, three dimensional model of the terrain. The model covers not only the heights but also models of how the surface elements reflect incident light. Building this inference system involves the development of a number of new technologies.

The inference performed about the terrain is the estimation of the parameters of a probability distribution over the surface. To enable further inference of the parameters of the terrain model when new images are acquired requires that we store sufficient information about the whole distribution to enable it to be used as the prior for the next update. The updates are performed either using a Kalman Filter, with sparse matrix representation of the covariance matrix, or by a gradient based optimization scheme on the posterior distribution, with re-estimation of the covariance matrix once the optimization is complete. An important consideration is how to prevent this matrix from becoming full as more observations are made.

The key inputs to the inference scheme are the real, observed image, the simulated image that the sensor would have recorded had it been observing the model rather than the real terrain, and a matrix of differentials, showing how the image would change as the parameters of the terrain model change. The technology of producing simulated images from terrain models is well studied in computer graphics. The technology of producing the differentials given a terrain model and a sensor model has, to our knowledge, never been studied before. A program which computes these derivatives for certain surface representation types and image formation models has been completed, and the underlying theory has been shown to be easily generalizable to other representations and models.

3-D Super Resolution from Multiple Images

Involvement in the ELVIS 3-dimensional super resolution project has been on two fronts. The design, implementation and documentation of a scientifically accurate renderer which is the forward problem for this project. Ensuring the accuracy of the derivations of the Generalized Kalman filter, which will be used to perform the inversion of the image data into the 3d surface model.

Graduate Student on Mars Project

Participation in the Graduate Student on Mars project has mainly concerned the classification of multi-spectral image data from the IMP (Imager for Mars Pathfinder) camera, involving pre-processing of the images from the left and right cameras before classification, and the production of a graphical utility to allow the domain specialists to easily visualize and interpret the results of the classification.

August 14th - September 30th was spent at the Signal Processing Laboratory, Cambridge University, working on Perfect Simulation using Markov Chains.

Input on the Graduate Student on Mars project has involved the pre-processing and post-processing of image and classification data. The images returned by the IMP (Imager for Mars Pathfinder) come from two cameras, corresponding to the left and right eyes. The two eyes record data in different spectral bands. To utilize all the data for automatic classification the images from the two eyes was first co-registered. The multi-spectral data was then passed to Auto-class, which produced an unsupervised classification. A graphical tool was produced to enable the domain scientists (geologists) to easily visualize the classes produced to determine the geological significance of the classes.

AGENT ARCHITECTURES

Barney Pell, Ph.D. (RIACS)

Investigated the design of Agent Architectures to enable highly autonomous operations of robotic and software control systems operating in remote, uncertain, and extreme conditions. Examples of such autonomous systems are robotic space probes, planetary rovers, antarctic submarines, and smart unmanned airplanes. These systems must be able to plan courses of activities to achieve high-level goals, schedule those activities to ensure that their plans fall within resource limits, and then execute the plans and track their internal health and the external environment. The agents must be adaptive in the presence of new information, reason under uncertainty to infer the most likely state of the world, and revise their plans on the basis of new information while maintaining a high level of reliability and responsiveness throughout. An agent architecture specifies how these capabilities must be designed so that they can be integrated into a complete system to achieve a diverse set of engineering design goals (reliability, robustness, maintainability, reuse, testability, etc).

The approach to research is applications-driven. Preference is given to challenging applications that push the state of the art. Real applications in particular to help distinguish those theoretical issues that must be practically addressed from those, which can be worked-around. They also raise challenges we have not considered, and frequently issues, which seemed to be mere details of implementation, integration, operations, or performance, turn out through practice to lead to the next generation of theoretical challenges.

Currently, RIACS is working with its own research team on four major application areas for autonomous agents: Mission operations, management of complex scientific instruments, communication satellite networks, and mobile robotics,

In addition to those applied research themes, RIACS is conducting a more theoretical investigation of some of the issues, which have emerged over the past few years of our applied research. These themes include the following:

1. Human-centered autonomy: The goal is to develop an agent which can minimize the necessity for human monitoring and guidance, but also maximize the opportunity for human interaction. This involves supporting multiple levels of commanding and monitoring across different activities and mission phases.
2. Real-time guarantees: To what extent can we guarantee that our autonomous agents will achieve their goals, relative to given assumptions about the world and real-time constraints? If the agent takes unbounded time thinking, it is hard to be sure that it can accomplish its tasks. One approach we are exploring is generating plans based on a model of the executive control system that will use the plan. The executive provides bounded computation guarantees, and if the plan respects those bounds then this provides provable knowledge. A related issue involves planning for contingencies, so that even under the presence of failures the plan can be guaranteed to succeed.
3. Agent capabilities: Is there a set of atomic capabilities of an autonomous agent that could be factored out into well-defined components? If so, these components and their interfaces could be standardized, supporting a more open architecture that could be built upon by the research community.

REMOTE AGENT EXPERIENCE

Ari Jonsson, Ph.D. (RIACS)

During the fiscal year of 1998, the main focus of research was within the field of automated reasoning, in particular constraint satisfaction, planning and scheduling. This was done in response to the need for reliable, autonomous systems, both for ground support and for flight software. Such systems are a key element of the move towards faster, better and cheaper missions, as autonomous model-based systems require fewer support personnel, are easier to adapt to new missions and improve science return.

Much of the work in this period was centered on the technology developed for and used by the Remote Agent Experiment, which is a part of the Deep Space One mission. The Remote Agent is a fully capable autonomous flight system that will be used to run the spacecraft autonomously for a period of six days. A central part of the Remote Agent system is an autonomous planner that can build detailed schedules of future actions, based on a set of high level goal requests and the current state of the spacecraft, while taking into account domain constraints, resource limitations, flight rules, etc. This planner consists of three layers, the topmost one being a search component that makes planning decisions, in an effort to find a valid plan. The next

layer applies planning decisions to a set of timelines, each of which describes the evolution of the state of an entity or system component. The consequences of the information on the timelines and in the model are eventually translated into the bottom layer, which is a dynamic constraint network. When all open planning decisions have been made such that the constraint network has a solution, the system has generated a valid plan. This planning system has been and continues to be the focus of our research work.

Research work in the fiscal year of 1998 can roughly be divided into five parts:

1. Work on the Remote Agent planner itself
2. Application of current planner to other domains
3. Design and development of the next generation planner
4. The design and development of new framework for better and more efficient dynamic constraint network management
5. Planner-independent research into reducing planning complexity by deriving useful relations from domain descriptions

Remote Agent Experiment

As noted above, the RAX planner separates the planning decisions from the reasoning about the lower level timelines. The management of the underlying timelines is done by the Heuristic Scheduling Testbed System (HSTS), while a plan search engine is used to choose among possible planning decisions. The search engine in the RAX planner is called "Incremental Refinement Search" or IRS. During the last fiscal year, our work on the search engine was twofold; to determine and document the engine's inner workings and to use this data to speed up the search process.

Our analysis of the search engine resulted in a detailed documentation of its functionality, its interface to the underlying HSTS system and its performance. Based on this, we identified a number of techniques that could be utilized to speed up the search process. This was an important task, as the search engine consumes a significant portion of the overall planning time. We discovered that two relatively minor changes, to simplify the decision choice point selection and eliminate provably unnecessary choice points, improved the efficiency by 10-15%. These improvements have been incorporated into the current RAX planner. Our analysis also led to other improvements, such as simpler and more efficient heuristic evaluation techniques.

Applying RA Planner to Other Domains

An important aspect of developing and providing general autonomous systems is to show that the methods and systems are indeed generally applicable. And even completely general frameworks must be shown to be general in practice, preferably by testing them on problems that other people have invented. To do this for our planning system, we applied the RA planner to a hypothetical spacecraft-planning scenario, developed by members of the TMOD group at JPL. The task at hand was to handle the science and communication planning for a Cassini-like spacecraft.

RIACS generated a new domain model that describes the spacecraft systems in question. Those included the main science and communication systems, the relation to Deep Space Network communications windows, the power available, the storage capacity, and many more. We then extended this model to include reasonably simple heuristic information about preferences between different options. This proved to be sufficient to allow the system to plan for a weeklong scenario in less than ten minutes. The end result was that the planner is more than general enough to handle this scenario quite effectively.

The planner and the model were also demonstrated to members of the TMOD group at JPL, including the authors of the scenario. Their evaluation was that the model was an accurate representation of the hypothetical spacecraft systems and that the planner had been shown to be capable of dealing with the given scenario.

During this effort, we also identified a number of problems with the current RAX planner. These ranged from serious theoretical problems to simple implementation mistakes. As a result of the discovery and analysis of these problems, the RAX planner a more reliable system.

Next generation planner

The current focus of our research is the design and development of the next generation of the Remote Agent autonomous planning system. This effort is aimed at developing a well-defined planning framework, based on the successful components of the RAX planner. This new planner will incorporate a number of new techniques and results that we expect to make the planner easier to use, more capable and more efficient. The new system is not only a new implementation, but a complete redesign that is based on clear concept definitions, well-defined interfaces and independent modules.

As part of this effort, during the last fiscal year, we identified a number of technical and semantical issues with the current system. We resolved most of these issues, which resulted in a simpler and yet more powerful definition of our planning paradigm. Based on this, we then developed an overall design for the new system. We also started on the component design and the actual implementation.

Improved Constraint Network Manager

One of the key technologies that we have incorporated into the new planning system, is a new general constraint network management system. We developed this system in an effort to make the core constraint handling of the planning system more general and more efficient. The result of that work was a general-purpose dynamic constraint network management framework that allows us to extend most existing constraint handling approaches to utilize any number of more efficient procedural methods, without losing useful properties such as correctness and consistency guarantees.

Our approach builds on recent results in the field of constraint satisfaction, which provided a theoretical foundation for the use of procedural methods with general search engines. This work was based on viewing a procedure as a program that can assign values to uninstantiated variables in a static constraint problem. We extended these results to dynamic constraint networks, by viewing procedures as external programs that eliminate values from variable domains, rather than assign values to variables. This generalization made it possible to utilize procedures in dynamic constraint networks, as procedural constraints. However, since some of the earlier theoretical results did not translate directly to this generalization, we had to establish new conditions that sufficed to guarantee desirable qualities like the maintenance of certain levels of consistency and overall correctness.

This new framework makes it possible to use temporal variables and temporal constraints, such as those required to do the temporal reasoning in the HSTS system, in conjunction with any set of general variables and procedural constraints. This is an essential quality for the new planner design, as there exist special-purpose methods for effectively reasoning about the temporal variables and constraints. Since the framework can incorporate any reasoning techniques, the result is a constraint networks manager that can effectively reason about temporal constraints, while retaining full generality and the capability of handling procedural methods.

As part of the development of the new planning system, we completed the design for the set of modules that implement the framework, then incorporated those into the overall design and tested the design by building a functional prototype.

Other planning research

An important aspect of the development of our planning systems is how they relate to other planning paradigms and systems. This gives rise to such questions as how the expressiveness of HSTS planning compares to the expressiveness of standard approaches such as STRIPS planning. To partially answer the question of relating HSTS to classical STRIPS, we developed a formal description of how propositional STRIPS problems can be mapped into temporal constraint problems. Since HSTS is capable of handling the temporal constraint problems, we showed that it is fully capable of dealing with propositional STRIPS problems.

This work then led us to another research area, the use of reasoning to eliminate unnecessary possibilities from the planning effort. The motivation is that for it is often easy for humans to look at planning problems and identify relations among states and actions that can significantly simplify the planning process. However, for computers to derive this information is often as expensive as solving the original planning problem. We made some significant progress in addressing this problem by developing the concept of restricted state spaces. The basic idea is to generate the state space for each type of assertion, independent of other types. This minimizes, and often eliminates; the exponential growth that is usually encountered in general state spaces. Actions are then added to show transitions between states, and they are linked to the set of states that allow the action to take place, i.e, the preconditions. The result is a restricted state space for the planning problem. Using this concept, we then developed automatic methods for identifying necessary states, equivalent states and dominated paths. This is an important result as the set of necessary states can be used to better guide the search effort, the equivalent states can be collapsed into a single state and any dominated paths can be eliminated from consideration.

BAYESIAN INFERENCE

Peter Cheeseman, PhD. (RIACS)

Research activity for FY98 focused on two main projects, as outlined below:

- 1) 3-D Super-Resolution from Multiple Images [in conjunction with Robin Morris, and other Caelum employees]

Following the successes of the earlier 2-D super-resolution system, as evidenced in the PathFinder Super-Resolution images, research commenced in earnest on extending the Bayesian approach to 3-D super resolution. This research is a considerable advance on the previous work as we were able, to extend the model to not only represent a 3-D surface, but to represent the uncertainty of the model at each location. This uncertainty representation allows us to recursively update the model as new images are processed. This research is being implemented using a 3-D triangulated mesh, but more recent research shows considerable advantages in using a smooth interpolated grid. The main achievements of this research are the detailed working out of the model update procedure and model representation, and the development of efficient algorithms for performing the model update. The success (or not) of these developments will become apparent during the implementation phase that was begun during FY98.

2) On Board Science Data Processing [joint research with Code S].

The goal of this project is to use on board processing of data (mainly images) collected by a planetary rover to decide what data or data summaries should be downlinked to earth. This is important because of the limited bandwidth usually available for downlinking compared to the amount of data collected. This research produced interesting science results by classifying [using AutoClass] the 12 band image pixels from Mars PathFinder. These results revealed that the ambient light is significantly spectrally shifted (due to dust scattering) so that shadowed areas (largely illuminated by ambient light) look spectrally different from non shadowed areas. This effect of viewing geometry on observed spectral distribution showed that simple computer vision algorithms are not sufficient for doing classification based on spectra, because they do not take the observational geometry into account. Owing to difference in opinion on the future direction of this project, Code I (i.e. P.C. and associates) withdrew from this project in August 1998, so that it is now an entirely Code S project.

B. HUMAN CENTER COMPUTING

SPEECH UNDERSTANDING

Frankie James, PhD. (RIACS)

In order to establish a baseline metric for future research in grammatical (and other) models for speech understanding, the project team has been developing statistical language models that will use trigram probabilities to predict spoken text. Currently, the focus is on calculating the trigram probabilities for a training corpus, which can then be used for prediction. The results of the prediction phase will be compared against future models.

Several trigram models are in the process of being implemented. All of these models use linear interpolation between raw trigram counts and different backing-off models. The trigram models differ in their handling of "special" cases, such as trigrams that only appear once, and words that may appear more than once but that always appear in the same context. The backing-off models differ in similar ways.

BRAHMS MULTI-AGENT MODELING LANGUAGE

Maarten Sierhuis, Ph.D. (RIACS)

RIACS has focused on the development and use of the Brahms multi-agent modeling and simulation environment. Brahms implements a new modeling paradigm for simulating work practices of organizations. The purpose of this research is to develop a new systems design paradigm for software systems that incorporates the understanding of the users' context of work. A common problem in the development of software systems is its acceptance by the end-user. Most software systems have a major impact on how people perform their work. User acceptance of a system is often based on the way the system can be integrated into the user's current work practice. In order to develop software systems that seamlessly integrate with a user's work practice; software developers need to understand the work practices of the user community.

RIACS is currently leading the Brahms project to develop a modeling and simulation environment (Brahms) that allows us to model and simulate the work practice of a group of people in an organization. Brahms consists of a number of different components:

1. A multi-agent modeling language. This is a newly designed language, based on other known agent-languages like Agent-0, and PLACA. The Brahms language allows us to develop agents that model the in-situ behavior of people in their context of work, including communication, movement, use of artifacts, etc. Each Brahms agent is like an independent knowledge-based system that collaborates with other agents.
2. A simulation engine. This newly designed engine can execute a Brahms model to create a dynamic model of the work practice of the agents in the model. The Brahms simulation engine is a discrete-event simulator.
3. Graphical end-user interfaces. Currently, we are developing two end-user interfaces that can show the result of a Brahms simulation: an agent activity time-line view, and a workflow view. Both views allow the end-user to analyze the results of a work practice simulation.

Our first objective is to do research on the application of the Brahms modeling approach within NASA center projects, with the primary goal to understand how to model work practices within NASA centers. The second goal is to understand the usefulness of such models in the context of software and organizational design and development projects.

UNDERSTANDING WORK PRACTICES FOR COMPUTER SYSTEMS DESIGN

Erik Vinkhuyzen, Ph.D. (RIACS)

Social scientific field research methods form a cornerstone of human-centered computing. Valerie Shalin from Wright State university and we started an extensive field research project with the Orbit Flight Dynamics Group in JSC this spring in order to study their work practices when on console in the Mission Control Center's Front Control Room. We applied several different methods, including observation, interviews and video recordings.

Some of the research topics that will be addressed in this project include:

1. the way in which the Flight Dynamics Officers coordinate their work in the Front Control Room with their colleagues in the front control room, as well as their support in the backroom, using the voice loops system, as well as other methods of communication;
2. the design of the computer tools the Flight Dynamics Officers use during their work. This research will consider the interface, or displays of the computers, as well as the underlying architecture
3. the way Flight Dynamics Officers are currently trained in Mission Control, and possible improvements to their training.

RIACS researchers have also started to collaborate with Nancy Smith other members of the CTAS team to help with the study of TRACON Air Traffic Controllers' work practices. Focussing on their specific strategies for separating descending aircraft.

3-D IMAGE RECONSTRUCTION

Xander Twombly, Ph.D. (RIACS)

A major focus of the NASA Ames Center for Bioinformatics is the generation of high fidelity 3D surface models from intensity mapped serial section data. The models consist of a contiguous mesh, which wraps around the exterior of an object, allowing both delineation from the surrounding objects and providing a model, which may be manipulated within the scene. Our goal is to provide both a computational environment and virtual tools, which allow the reconstruction and manipulation of such data sets for use in scientific and medical investigations.

Anatomical reconstruction of portions of the human body from CT or MRI data sets represent one area of research being actively pursued in conjunction with the academic and medical communities. From a computational standpoint, the high fidelity images generated from CT and MRI data of the human body yield a number of interesting problems. A typical reconstruction of the muscle and soft tissue boundaries of the heart might require a mesh containing 800,000 triangles, while a human skull might contain upwards of 2,000,000 triangles. Even on today's graphic supercomputers such as the Onyx2 Infinite Reality Monster, manipulation of these large objects is quite difficult and tasks as simple as rendering and displaying the objects can take upwards of a second to accomplish. On less expensive (and less powerful) equipment, rendering of these scenes can take in the tens of seconds to perform. These slow rendering rates restricts the utility of these images to simple displays of the data, and tend to preclude the use of an interactive environment due to the long rendering times. Our goal is to design a system that remains faithful to using the high fidelity images obtained from the original anatomical data, allows real time interaction with the data for manipulating and modifying (i.e. cutting) the mesh, provides a mechanism for multiple users to work on the data simultaneously, and allows access from low end workstations without sacrificing the rendering speed of a graphics supercomputer.

Virtual Distributed Online Clinic (vDOC)

The virtual Distributed Online Clinic is an effort to provide a group of practitioners with the ability to interact with both the patient specific 3D reconstruction and each other during diagnosis and treatment planning. A collaborative, interactive environment for the 3D data can be produced in two general ways: 1) Independent workstations that manipulate and display the data locally, and transmit the changes enacted locally to all the other workstation members of the session, or 2) A central server designed to perform all of the data manipulation and display, and a scheme to replicate the server display to all the clients in the session. Each approach requires the use of expensive, high end computing equipment, but they can be differentiated primarily by their computational and network requirements.

The first technique has the advantage of a fairly low bandwidth requirement on the network. Once the initial 3D-mesh data set has been distributed to the client workstations, the information transmitted between sites is fairly minimal. Information such as rotations, markings (with a pointer device), and even cuts in the mesh structure can be transmitted using a minimal amount of bandwidth, making such a system feasible over today's internet connections. The drawback is that each workstation must have the hardware capabilities to support high fidelity models, including the display speeds required for interactive manipulation. For large data sets such as those generated from anatomical images, this virtually requires a graphics supercomputer to be on site, thus limiting the use of such a collaborative system to a select few that would access to these computational facilities.

The second approach is to trade the power of the client workstations for the bandwidth of the network connections. If all computation and display is performed by a central server, then the client workstations can function simply as a remote display, and do not require any special hardware. Such a system would replicate the display seen on the server to each of the remote sites, and any manipulation performed remotely would be transmitted and executed by the central server. The drawback to this design is that it requires a tremendous amount of bandwidth relative to today's networks, and could potentially saturate a system as large as an ATM line (nominal bandwidth of 155Mbps).

Since July 1998, our system has been modeled on the client/server design, with a focus on efficient transmission of the images across the network. Initially, our goal was to replicate a display that was 1024x1024 pixels in size, with 24-bit color (TRUE COLOR) depth for each pixel. Ideally, when an object is in motion (e.g., rotating) it will be redisplayed 15 times per second, to give the illusion of smooth animated motion. Finally, the objects are viewed in visual stereo, so we would need to provide an independent view for each eye from any given vantage point. Taken together, these specifications provided a daunting bandwidth requirement:

$$2(\text{stereo}) * 1024 * 1024(\text{pixels/frame}) * 24(\text{bits/pixel}) * 15(\text{frames/second}) = 560\text{Mb/sec.}$$

Using these specifications as a starting guideline, the vDOC system has evolved as follows:

- 1) An OC-12 ATM network can provide 622Mb/sec bandwidth. However, very few interface cards exist that function at this speed on an SGI Onyx2 Reality Monster, which is the host graphics supercomputer. Furthermore, no OC-12 network cards exist for low-end SGI workstations and PC desktop machines, so the client computers would be unable to interface to the network at these speeds. Consequently, an OC-3 ATM connection with a putative 155Mb/s transmission (max of 132 Mb/s is to be expected) capability is the greatest bandwidth that we can expect to use. Current wide distribution of fast Ethernet (spec'ed at 100Mb/s throughput, max expected throughput of 70-80Mb/s) also leads to the question of compressing the signal to make use of current network infrastructure.
- 2) Signal compression is being performed in 3 ways. The first is to employ a lower color resolution, using only 5 bits per color instead of 8 bits per color. This allows the compression of each pixel into a single 16bit value, while retaining a large degree of color fidelity. Investigation into the use of color maps is also underway to try and maximize the color range. The second form of compression is to implement a run length encoding (RLE) scheme. Compression varies depending on image content, but tends to produce a compression ratio of 2-4:1. Finally, stereo compression is performed by generating the stereo images in an under/over format, which simply compresses the horizontal resolution of each image by a factor of 2 so that views for both eyes can fit on a single frame. Taken together, these compression schemes reduce the required network bandwidth to 60-120 Mb/s, which falls well within the capacity of a dedicated ATM line, and possibly within the capacity of a fast Ethernet connection.
- 3) Dynamic resolution is an aspect of the problem that is currently under investigation. Our premise is that when an object is in motion, the view does not require the level of spatial detail in the image normally used when the object is static. Dynamically decreasing the level of detail while the user is performing certain operations on the data (such as the rotation of an object) can enhance the illusion of interactivity significantly while concurrently reducing the required bandwidth. Current experiments with reducing the display to half the vertical and horizontal resolutions shows very little reduction in utility of the software, and reduces the bandwidth requirements by 75%. The advantage to this approach is reducing the bandwidth to 15-30Mb/s, which can readily be provided by the fast Ethernet infrastructure.

- 4) Hybrid client/server designs are also under investigation. This is a modification of the basic client/client server design that allows the client to act as a workstation under some circumstances, increasing response time by performing calculations on the local machine while objects are in motion. In this scenario the decrease in resolution used during movement is not a reduction in image size, but rather a reduction in detail of the mesh. Thus a model of the heart normally viewed at 800,000 triangles might be rendered locally using only 50,000 triangles. Current rendering speeds on a mid-range PC running Windows NT can accommodate rendering objects of this size at near real time speeds. While the object is in motion, the local hardware would render and display the scene and then request the final static frame (at full resolution) from the server. The advantage of this system is both the reduction in required bandwidth and the ability to take advantage of whatever hardware is available at a particular client site.

- 5) Multipipe rendering on the Onyx2 Infinite Reality Monster is also under investigation for the largest data sets. The maximum rendering speed of a single graphics subsystem on an Onyx2 is 10M triangles/sec. Realistically, the rendering speeds are only about 2 - 4M triangles/sec, which is an order of magnitude below what is required for displaying the human skull at 15 fps. Working with SGI, we are developing ways to use multiple independent graphics subsystems on the same Infinite Reality Monster to interleave the frame generation. This is a fledgling technology on SGI's part, and we are actively developing ways to use their newest designs to encapsulate the multipipe architecture of the Onyx2.

The vDOC system is still in the early development stage. We have committed to a demonstration of this technology in January 1999, linking client sites at Ames, Stanford Hospital, Cleveland Clinic, Salinas Valley Memorial Hospital, and the Navajo Nation in New Mexico, all driven from a graphics server at AMES.

CT, MRA AND 3D TIME VARYING ULTRASOUND DATA

Steven Senger, Ph.D. (RIACS)

RIACS spent time in the development of techniques to preprocess CT, MRI and 3D time varying ultrasound data to facilitate the segmentation of soft tissue and the construction of surface models. The work is focused primarily on the heart because of its importance to NASA space flight research. All of the imaging modalities produce data sets, which exhibit noise or artifacts of one form or another. CT typically contains radial shadows produced by high-density material. MRI typically exhibits signal falloff across the imaged region. 3D ultrasound is a research technology, which provides relatively coarse data. Standard image processing algorithms are capable of reducing the effect of the noise but tend to remove fine detail such as the surface blood vessels of the heart. On this project we have continued the development of techniques which produce clean surface models but which also retain fine detail in the original data set.

The development of a NGI application to provide interactive access to high performance computer graphics to geographically distributed sites. This project, called vDOC (Virtual Distributed Online Clinic) will allow physicians to collaboratively visualize and manipulate 3D reconstruction of patient specific data. A central imaging server will provide remote client sites with high quality, interactive images of surface models constructed from patient specific data. The sites will be connected through NASA NREN and satellite links. On this project we have consulted on the design and implementation of the visualization and network transport software.

The development of surface models of the planet mars which incorporate altimeter data and surface geology. The goal is to develop a surface model, which not only displays the topography of the planet but also incorporates best knowledge of the visible surface geology. On this project we have consulted on computer graphics techniques for building and visualizing the desired surface models.

NEW VISITATION PROGRAM (ROSS6.1) WITH OPENGL

Rei Cheng, Ph.D. (RIACS)

Implemented new visualization program (ross6.1) with OpenGL, and completed the development of the new features onto our new visualization software (ross6.1). Such as, applying different colors for each object's interior color and exterior color. Allowing users to define their own arbitrary clipping plane, and to define the position of light. As well as getting the frame rate while rendering in second window. To have some editing tools to modify the existing animation sequence; the menu item called "Modify Animation Sequence".

The functionalities include, modifying number of steps, keyframes, and display mode. To continue to fix new bugs found in this new version of ross6.1 program.

Integrated particular functions for special application we must simplify baby skull mesh for use in "SkullCut" program Integrated code for generating fairly simple mesh of baby skull for use in the program called "skullcut" - an interactive virtual Surgery program developed by Aaron Lee.

Evaluating the quality of multi-resolution modeling techniques Created an animation of 3-D reconstruction human heart with three different resolution meshes as follows:

- full resolution of heart (total 860k polygons)
- 1/17 resolution of heart (total 50k polygons)
- 1/86 resolution of heart (total 10k polygons)

The result shows two things. 1) how an efficient simplification algorithm (QSLim) can reproduce same high quality images with much less (1/10) meshes, and 2) how much the rendering time can be speed up with a much lower number of mesh size.

QSLim, is the software we use to simplify our high-density medical images (generated from 3-D reconstruction of CAT scan or MRI)? It was developed by Michael Garland; this multi-resolution modeling technique based on the algorithm called "Edge Contraction".

Extended this simplification algorithm to very complicated human face. Another major effort was to reconstruct the visible human's face in four different level details of meshes. The file information as follows:

- full resolution of face (total one million polygons)
- half resolution of face (total 500k polygons)
- 1/4 resolution of face (total 250k polygons)
- 1/10 resolution of face (total 100k polygons)

The 3-D reconstruction of her face and skull had been done successfully. The new animation sequence generated from this new data set shows how an efficient algorithm can lead to the extent of simplification process. Especially, from interior of the skull you can see clearly how the fissure of her brain press the skull.

Technical support was provided for our summer visiting fellowships and students from universities and outside research laboratories as well. The research projects include 3-D Computer Reconstruction of the human brain and skull (Ratan Bhardwaj, Medical student from Queen's Univ., Kingston, Ontario, and Canada). Virtual Reality Simulation of the effects of micro-gravity in gastrointestinal physiology (Cesar Compadre, professor from Depts. of Biopharmaceutical Science, Univ. of Arkansas for Medical Science). 3-D reconstruction of the visible female's pelvis bone, hands, feet's, and others. (Yasmeen Nkrumah, college student). Interactive virtual Surgery program (a innovative cutting tool) (Implemented by Aaron Lee, Ph.D. student from Univ. of Princeton)

Designed and provided new images for our "Center for Bioinformatics". Generated image files from these new medical data sets and put together to make new posters for task requester to display at Salinas Memorial Hospital Open House and at VH '98 exhibit. Designed four key composite images for our Virtual Hospital project, and generated some hard copies and slides for Dr. Ross to present in very important conferences and for various purpose as well.

Produced new Video productions. RIACS had generated a couple different versions of animation video taped for task requester to show in the national conference, to provide the media (TV stations and newspapers) and also to be presented to our visitors from time to time in our center as well. Presented Immersive Workbench demo with stereoscopic view in numerous our center's demo. Major efforts were put together each data in certain specific way then generated animation with each data's unique quality. Those demos are always very impressive to all visitors and have been praised by the task requester. It includes demo_baby, demo_lung, demo_heart, demo_breast tumor, demo_teenboy, demo_girl

3D ULTRASOUND HEART MOTION SEQUENCE

Esfandiar Bandari, Ph.D. (RIACS)

The work so far has included getting a video camera for the Electron Microscopy Lab functional and working on two projects: 1) volume reconstruction from 3D ultrasound heart motion sequence, and 2) 3D reconstruction of the Mars surface from MGS (Mars Global Surveyor) and Viking data. The following will show some preliminary results from both projects. These researches were conducted in Matlab on a PC.

Volume Reconstruction from 3D Ultrasound Heart Motion Sequence: The 3D-ultrasound sequence that has been provided to us suffers from quality problems. That is, noise and constructive as well as destructive speckle seem to have corrupted the reflections from the heart muscle itself. Due to the large data set, only a portion of the data was extracted and analyzed. Dr. Steve Senger had originally worked on this data set using motion analysis to track the volumetric data in hope of reconstructing a 3D set by back projecting the sequence in time, averaging the data and reducing the noise.

RIACS approach to the problem is slightly different. Instead of tracking the raw data, we decided to preprocess the volumetric data to reduce the effects of noise and speckle. RIACS would like to identify areas of interest in the image for further processing later. This basically amounts to the removal of the void space often filled with noise and speckle artifacts from further consideration. Another benefit of direct 3D reconstruction is that it allows its application to non-moving parts of human anatomy such as liver, kidneys and bladder among other areas.

The figure below shows a typical example of the raw ultrasound image. The histogram of these cross sections shows the typical Rayleigh (and higher order Rician) distributions. As we will see later this becomes an important fact in the sequence of preprocessing that will take effect.




Figure 1. An Example of the 3D ultrasound image slice

Several examples of image restoration for ultrasound include the use of homomorphic filtering to remove the image blurring artifacts.¹ This technique does require 3D phase unwrapping and the separation of the blurring function and the signal in the cepstral domain for clean removal of the blur. Since, the image suffers from other artifacts, not to mention that 3D phase unwrapping is quite unstable, the use of this approach was delayed. Another method considered for image segmentation was the use of Markov random fields and simulated annealing.² Looking at the results for 2D images and the time of convergence of this approach convinces one that similar or better results can be achieved by the use of stack filters.

Jonathan Carr's doctoral thesis³ on 3D ultrasound provided some interesting insights. Even though the work presented no real techniques for noise and speckle removal, it did provide references and discussions of various filters, which were used in this project.

To remove the effects of the long tail Rician distribution, a 3x3x3 morphological erosion filter (which is the same as a min filter) was first used on the image. This filter has the additional interesting effect that it removes noise and speckle with large values. Additional noise and speckle was then removed by the use of a 3x3x3 median filter. The result shifts the distribution of the pixel values more toward the middle, thus making the signal having a more Gaussian distribution. Finally to fill the gap left by the last two filters, a 3x3x3 morphological expansion operator was used. This amounts to the use of a max filter. The result of these operations is displayed in figure 2.

¹ Torfinn Taxt, "Reconstruction of Medical Ultrasound Images Using Two-Dimensional Homomorphic Deconvolution", IEEE Trans. Ultrasonics, Ferroelectrics and Frequency Control, Vol 42, No. 4, pp 534-554 July 1995

² Stewart Geman, and Donald Geman, "Stochastic Relaxation, Gibbs Distribution and the Bayesian Restoration of Images", IEEE Trans. Pattern Analysis and Machine Intelligence, Vol 6, No 6, pp 721-741 Nov 1984.

³ Jonathon Carr, "Surface Reconstruction in 3D Medical Imaging", Doctoral Dissertation, Department of Electrical and Electronic Engineering, University of Canterbury, Christchurch, New Zealand Feb 1997.




Figure 2. Segmentation of the heart image after the preprocessing

Figure 2 shows the preliminary results generated with the above technique. As it can be seen many of the speckle and noise effects in the chest and heart cavity have been removed and better structure of the heart volume is displayed.

Problem with the Present Approach: The main problem with the present approach is that even though it provides a good mask for regions of interest in the image, the result is too blocky. One remedy which is easy to implement, but has not been tried yet is to use a more circular filter rather than a 3x3x3 filters. This should give better results but can also cause a bias toward orientation of the filter for such a small region. A remedy can be to use a larger and more circular window for the median filtering. Yet another approach can be to binarize the above volumetric image, and use a small circular mask to "open" the image morphologically.

Other Results and Conclusions: Using a 3D Laplacian of Gaussian filter three-dimensional planar streaks were discovered. These capture artifacts can cause serious problems in tracking since they will force the matcher to lock into periodic. Further work in this area will continue in near future.

3D Reconstruction of Mars Surface From Satellite Imagery: In order to expand the utility of the ROSS modeling software, the Center for Bioinformatic at NASA has been looking at 3D reconstruction of Mars surface from Viking and the Mars Global Surveyor (MGS) data. This work is particularly important given the availability of the MGS altimeter data and the significance of the high-resolution satellite imagery from MGS.

Our final approach will hopefully be one that will extend this work to motion stereo analysis for medical imagery as well as planetary surface reconstruction and obstacle avoidance for future missions to Mars. But for now we are working with single satellite imagery and using Shape from Shading approaches to reconstruct the surface of the red planet. The figure below shows a typical satellite image.




Figure 3. Image Mars from Mars Global Surveyor

Using localized Shape from Shading algorithms⁴ we have been able to get an approximation to the 3D structure of Mars; unfortunately not knowing the exact direction of light this is just an approximation but the main peaks and the valleys of the landscape are well displayed in the resulting image.




Figure 4. Result of Localized Shape From Shading.

⁴ Ping-Sing Tsai and Mubarak Shah, "A Fast Linear Shape from Shading", IEEE Conference on Computer Vision and Pattern Recognition, pp 734-736, 1992

As it can be seen these results are also preliminary, but given accurate direction of light, larger number of iterations and fusion of information from the MGS altimeter the results can be improved.

Future work will hopefully include, shape from motion stereo and multiple data fusions.

C. HIGH PERFORMANCE COMPUTING AND NETWORKS

NASA RESEARCH AND EDUCATION NETWORK/NEXT GENERATION INTERNET (NREN/NGI)

Marjory J. Johnson

The NREN/NGI project is NASA's contribution to the Federal NGI initiative. NREN is a next-generation network testbed connecting several NASA centers across the U.S. NREN is peering with NGI testbeds sponsored by other Federal agencies, with the CAIRN (Collaborative Advanced Interagency Research Network) testbed, and with the Internet2 testbed to provide a nationwide network testbed for conducting network research and demonstrating revolutionary applications.

M. Johnson is group lead for the applied research group within the NREN/NGI Project Office. M. Johnson is also a member of the Large Scale Networking - Networking Research Team (LSN-NRT), an interagency group whose overall objective is to facilitate collaboration on NGI research activities among the various Federal agencies.

The focus of the NREN project is to develop and demonstrate applications that will revolutionize the conduct of NASA science. To support development of these applications, NREN is conducting research primarily in the areas of multicast and Quality of Service (QoS).

QoS WORKSHOP

Quality of Service pertains to the ability to deliver a specified level of performance for a particular network application, regardless of contention for network resources from other applications.

Because QoS is critical to NASA missions, the NREN/NGI Project Office hosted a Quality of Service (QoS) workshop at NASA Ames Research Center on August 18 and 19, 1998. M. Johnson was co-organizer of the event. There were approximately 100 attendees, representing NREN partners from several NASA centers, other government agencies, universities, the Internet2 community, and industry.

In the past NASA has used dedicated circuits to support critical network applications, an expensive approach resulting in over provisioned networks that are idle much of the time. Emerging technological advances will enable more cost-efficient sharing of network resources by delivering QoS guarantees to selected user applications. That is, QoS will enable designated applications to receive preferential treatment at the expense of lower priority applications, when network resources are inadequate to satisfy all requests.

Presentations at the workshop were grouped into sessions based on key QoS topics. Each group of presentations was followed by a general discussion on that topic. Since QoS is meaningful only in the context of specific application requirements, the workshop opened with a discussion of two specific NASA applications: digital video distribution to support shuttle-launch operations at NASA Kennedy Space Center and Earth Observing System (EOS) data distribution. The bulk of the workshop was spent addressing specific QoS research areas, including QoS routing and middle ware, experimental measurement of QoS deliverables, plans for implementing a QoS infrastructure for Internet2, and QoS policy.

Copies of position papers and presentation slides are available on the NREN web site:

www.nren.nasa.gov

M. Johnson prepared a report, also available on the NREN web site, that highlights open research issues raised during the workshop.

NREN RESEARCH ACTIVITIES

Some NREN research activities will be conducted in-house; others will be conducted in partnership with other NASA centers, other Federal agencies, and the university community.

NREN is currently setting up a QoS testbed in the NREN lab. Joint activities are being planned and coordinated with NASA Goddard Space Flight Center and the National Institute of Standards and Technology (NIST). NREN is also planning collaborations with the CAIRN testbed and the Internet2 university project.

Another in-house research project, led by the Data Understanding group in Code IC at NASA Ames, is the application of AI technology to the problem of network routing. Code IC personnel are developing a learning methodology to dynamically load balance packet traffic within a network. NREN is providing simulation support and network analysis expertise to the project.

During the past year M. Johnson worked with several universities to develop research plans to support the NREN/NGI program. Grants and contracts have been awarded to five universities. The projects address issues of fault-tolerant multicast, allocation of network resources across network domains, performance and efficiency of IP over ATM versus IP over SONET for transmission of large volumes of data, development of a QoS Application Programming Interface, development of an architecture for a centralized policy-management system, and denial of service in the network infrastructure.

NREN/NGI APPLICATIONS

Candidate applications for NREN development and deployment must require network capabilities that are not provided by NASA's current operational networks, and they must promise to revolutionize the conduct of NASA science. NREN is currently pursuing activities in the following areas: support of shuttle-launch operations at NASA Kennedy Space Center, distribution of large data sets as part of the EOS mission, next generation aerospace design methodology and the Information Power Grid project, telemedicine, and development of a virtual enterprise in support of the Astrobiology Institute.

M. Johnson is responsible for the research component of NREN applications. Research will be required to enable revolutionary applications; further research will be required to transition these applications to the operational NASA Integrated Services Network (NISN).

M. Johnson is currently part of an NREN team that is helping to develop the virtual Distributed Online Clinic (vDOC) application, in conjunction with the Biocomputation Center at NASA Ames. The scientific objective of this application is to create a virtual hospital. Three-dimensional high-resolution images of the human body (obtained via X-rays, MRIs, or other data sets) will be transmitted across the network in near-real time. The vDOC application enables virtual surgery, by allowing an end user to rotate, manipulate, and slice into these images. It will help to enable the delivery of health care to remote areas of the world, and ultimately into space.

Many networking challenges must be addressed for successful deployment of the vDOC application. A demonstration of early capabilities is planned for January 1999; end-site participants will be NASA Ames Research Center, NASA Lewis Research Center/Cleveland Clinic, Stanford University Medical Center, Salinas Valley Memorial Hospital, and the Gallup Indian Medical Clinic.

NFR INTRUSION DETECTION SYSTEM

David Gehrt

The principle area of interest has been intrusion (and intrusion attempts) detection. We are currently conducting a test of the NFR intrusion detection system using N-code nodules from Anzen Computing. This started with log scans on the primary name server for the arc.nasa.gov zone which we operate for the Center, and which served as an early warning system for some types of attacks launched against systems at Ames.

RIACS was also part of the beta program for the RedHat 5.2 version of Linux on the Sun Sparc architecture, which is just concluding.

PARALLELIZATION OF IRREGULARLY STRUCTURED CODES

Leonid Oliker

The focus of our work was parallelization of irregularly structured codes. Our application has been unstructured adaptive grids, which are generally expected to constitute a significant fraction of future high-performance supercomputing. We have developed a framework called PLUM whose research issues include: dynamic load balancing, graph partitioning, data remapping, and communication cost modeling.

Our work is being well received in the scientific community, and during the past year we have produced a number of publication in distinguished journals and conferences. Highlights include: the acceptance in JPDC[7], the leading journal for parallel computing; a book chapter in the upcoming IMA Volumes[5]; and a publication in a leading aircraft journal[10]; Additionally, we presented a paper at Euro-Par98[6] in England this summer. This work was co-authored by Professor Harold Gabow, a leading graph theoretician. We are also collaborating with George Karypis and Vipin Kumar from the University of Minnesota. Their team has created MeTiS, possibly the most powerful partitioner in the world. Recently elements of PLUM were incorporated into MeTiS[4], meaning that future distributions of their software will include our technology.

It is our plan to continue developing optimization strategies for irregular computing in a distributed environment. Originally the mesh adaptation routine was written for the SP2 in a distributed memory fashion using MPI. Since NAS has acquired an Origin2000 we are in the process of rewriting the code using the shared memory CCNUMA paradigm[2]. A performance comparison between these two approaches will provide valuable insight for future code development. We also plan to examine various other programming strategies for performing adaptive computations. These include a Symmetric Broadcast Network (SBN)[1], multithreading techniques, and graph theoretic approaches such as self-avoiding walks.

PARALLEL PROGRAMMING

Susie Go

Parallel programming has emerged as an important tool in large-scale scientific computing. Although domain decomposition and multigrid methods lend themselves naturally to the parallel

environment, computing solutions over realistic but complicated geometry's also calls for the flexibility which unstructured grids offer. Difficulties occur when using unstructured grids with multilevel methods because a hierarchy of grids is needed but is not naturally available for unstructured grids. Our interest focuses on the analysis and construction of robust multigrid and domain decomposition methods applied to elliptic problems on general unstructured grids.

Previous analysis and computational approaches have shown that in the case of model elliptic problems, when unstructured multigrid methods are constructed to satisfy certain approximation and stability requirements, multigrid convergence can be achieved. Proper treatment of the non-nested spaces and non-matching domain boundaries, which often arise in the coarsening of unstructured grids for multilevel methods, is necessary in the construction of methods, which retain optimal convergence rates. Our approach uses the definition of accurate and stable interpolations between levels.

RIACS has extended these tools and analysis for applications beyond the model elliptic problem. An unstructured elliptic eigenvalue solver was developed by adapting an existing structured multigrid method for solving differential eigenproblems on uniform grids [Brandt-McCormick-Ruge 1983]. Results showed that the symmetric eigenvalue problem on unstructured grids could be solved with optimal convergence using such a multigrid algorithm.

The solution of eigenvalue problems also occurs in one approach to solving the NP-complete graph-partitioning problem, which is an important component of parallel computing, particularly in domain decomposition. Many partitioning methods and associated sophisticated software packages have been developed recently including algorithms based on greedy, coordinate, inertial, multilevel spectral and graph bisection. There is usually an unavoidable tradeoff between quality and speed. Of these, spectral partitioning has the reputation of being the most time consuming, but it also consistently produces high-quality partitions which for some applications can offset its costliness with better convergence rates of iterative solvers.

Attempts at accelerating the computation of the Fiedler vector (the lowest nontrivial eigenvector of the associated graph Laplacian) by multilevel methods encounters the difficulty of adapting standard elliptic multilevel algorithms to solving the discrete graph Laplacian eigenvalue problem. Thus, for example, it is difficult to construct such algorithms that have convergence rate independent of the mesh size.

RIACS has shown that when the graph is a finite element mesh, an equivalence between the discrete and continuous Laplacian operator exists and we use this to adapt recently developed multilevel elliptic algorithms on unstructured grids for solving the graph-partitioning problem with a true multigrid convergence rate. The unstructured multigrid eigenvalue solver can then be used for efficient solution of the graph-partitioning problem using the spectral bisection method [Simon-Pothen-Liou 1990].

Though multigrid methods are currently used for computing flow solutions, their convergence properties are far from optimal. The same unstructured multilevel techniques we have been working on can be implemented and studied for the Euler equations to generate more efficient algorithms. Development of software for finite element discretizations (the Element Library for Fluid Flow, ELF, with Timothy Barth of NASA Ames) and integration of these discretization routines with a standard library of linear and non-linear solvers (the Portable, Extensible Toolkit for Scientific Computations, PETSc, developed at Argonne National Labs) has been completed. Currently, we use Newton's method with GMRES (preconditioned with ILU(0)) which are provided by PETSc to solve the system of equations generated by the ELF discretizations. Convergence studies with the ELF/PETSc software for 2-D Euler flow on unstructured grids will include multigrid and 2-level Schwarz preconditioners using tools similar to those which have been developed for elliptic problems.

HYDROGENATED AMORPHOUS SILICON (A-SI-H) AND NANOCRYSTALLINE SILICON (NC-SI-H)

Ramalingam, Ph.D. (RIACS)

Hydrogenated amorphous silicon (a-Si:H) and nanocrystalline silicon (nc-Si:H) have emerged during the past decade as cost-effective alternatives to crystalline silicon (c-Si) in electronic, optoelectronic, and photovoltaic applications ranging from manufacturing of high-efficiency solar cells to thin-film transistors for flat panel displays. Films of a-Si:H and nc-Si:H are produced usually by plasma-enhanced chemical vapor deposition (PECVD) from silane and hydrogen (SiH_4/H_2) containing discharges. Previous experimental research has led to commonly acceptable ranges of PECVD parameters to achieve silicon films of given structural quality; such deposition parameters include, e.g., the substrate temperature and the feed gas composition. Despite the extensive study, however, the surface chemistry that controls the structural quality and the properties of the deposited Si films is not understood and, thus, the reaction mechanisms and the gas-phase precursors that lead to deposition are still unknown.

Motivated by the above fundamental materials problems in silicon PECVD, the **objective** of our USRA-funded research is two-fold:

- detailed atomistic study of plasma-surface interactions in the PECVD of hydrogenated silicon, and
- direct atomistic simulation of the deposition process over a range of PECVD conditions.

The computational methodology and specific research tasks are summarized below. Addressing the research objectives has resulted in the publications listed in Section II. Conference, seminar, and colloquia presentations of this research (exclusively or partly) are listed in Section III.

Computational Methodology

Our atomic-scale studies are based mainly on molecular-dynamics (MD) simulations and are aided by molecular static's (MS) and Monte Carlo (MC) simulations. The software required for executing and analyzing these simulations has been developed at UCSB over the past three years. The description of interatomic interactions that we employ in modeling the interactions between radicals originating in the plasma and surfaces of Si substrates and films during PECVD of hydrogenated silicon is based on an extension of Tersoff's potential for Si to incorporate Si-H, H-H, and the corresponding three-body interactions. We have tested the accuracy of this classical empirical interatomic description systematically and found it to be very satisfactory in describing the structure and properties of radicals in SiH_4/H_2 plasmas, bulk a-Si:H phases, and H-terminated c-Si surfaces. Most recently, the predictions of the interatomic potential have been found to be very satisfactory in comparison with quantum mechanical density-functional theory (DFT) calculations of the energetics of surface chemical reactions on crystalline Si surfaces performed by Dr. S. P. Walch at NASA Ames. Coarser-time-scale modeling of the PECVD process will be based on hybrid off-lattice kinetic Monte Carlo (KMC)

simulations that are being developed at UCSB and require as input a surface chemical reaction database that is being developed based on the results of the above atomistic simulations.

Research tasks in the study of radical-surface interactions

These tasks are described in some detail below. They aim at the chemical characterization of various surfaces of interest with respect to chemically reactive radicals originating in the plasma, as well as the detailed analysis of the dynamics and energetics of such radical-surface interactions.

- Characterization of surface chemical reactivity with respect to each radical for
 - Pristine Si(001)-(2x1) surface
 - Hydrogen-terminated Si(001)-(2x1) surface
 - Surfaces of a-Si:H films at various hydrogen surface coverages based on MS methods of energy minimization. Chemical reactivity maps have been constructed for the SiH, SiH₂, and SiH₃ radicals. The amorphous films on crystalline substrate systems have been generated by a combination of MC, MS, and MD methods.
- Analysis of the interactions of the SiH, SiH₂, and SiH₃ radicals with the above surfaces based on MD simulations: the MD trajectories are calculated after impinging the radicals on the surface at a given molecular orientation and at surface locations on a grid constructed according to the chemical reactivity characterization of the surface. MD simulations have been completed for the SiH and SiH₃ radicals; simulations for the SiH₂ radical will be completed during the next fiscal year.
- Detailed analysis of surface chemical reactions identified by classical MD trajectories: this analysis involves MS quenching of MD trajectories and construction of classical reaction paths, as well as comparison of the corresponding energetics with DFT calculations. Some of this analysis has been completed for the SiH₃ radical, including the abstraction of hydrogen by SiH₃ from the H-terminated Si (001)-(2x1) surface, which is believed to be one of the key reactions during Si PECVD. The remaining analysis for SiH₃, as well as the analysis for the SiH and SiH₂ radicals will be completed during the next fiscal year.
- Development of reaction rate database: this is based on the above reaction path energetics combined with variational transition-state theory for rate computations. The database for the key reactions identified by the above analysis will be completed during the next fiscal year. This database will provide some of the necessary input for KMC modeling of the PECVD process.

The DFT calculations in the above analysis have been performed by S. P. Walch at NASA Ames. The MD trajectories are used to provide various initial configurations for the DFT calculations, thus minimizing the effort required for the computationally more intensive DFT calculations. In addition, comparisons between the MD and DFT results will validate further the capabilities of the classical interatomic potential for accurate reaction rate computations. It is

anticipated that this joint effort will be strengthened during the next fiscal year and lead to the efficient development of an accurate surface reaction database for PECVD process modeling.

The chemical reactions that are studied in this part of the research are isolated in the sense that a single radical interacts with the Si surface. Chemical reactions also are studied in the next part, i.e., during deposition; these reactions are not isolated events, since they occur simultaneously with relaxation of the surface on which species keep impinging. Comparison of both types of reaction paths will help validate certain assumptions in the implementation of the KMC coarse-time-scale simulations.

Research tasks in the modeling of Si plasma deposition

At the current stage of the research, plasma deposition is studied approximately and in an accelerated fashion by MD simulations assuming that:

- the SiH_3 radical is the only species impinging on the surface and
- the rate of radical impingement is determined by MD time scales, where a silyl radical is brought to the surface every few picoseconds.

The radicals impinge on an initially H-terminated surface normally and at random molecular orientation. Surface chemical reactions are classified into two groups: those that occur in the initial stage of deposition, where the radicals react with the original crystalline surface, and those that occur at later stages, where the radicals react with the surface of the deposited amorphous film. The range of the substrate temperature, T , over which the process is studied is $300 \text{ K} \leq T \leq 773 \text{ K}$, as implied by actual PECVD conditions except for the high end of this range that is used in the simulations to accelerate the thermally activated physicochemical processes that occur on the surface.

The main objective of these approximate nanosecond-time-scale deposition simulations is the identification, classification, analysis, and rate calculations of certain chemical reactions that play important roles in Si PECVD. Thus, the specific research tasks of this part of the study include:

- Generation of long MD trajectories under the above assumptions for further analysis. The substrate temperature ranges $500 \text{ K} \leq T \leq 773 \text{ K}$ has been investigated. Key reactions that have been identified during the initial stage of deposition include: (i) H abstraction by SiH_3 to release SiH_4 molecule into the gas phase with simultaneous formation of surface dangling bonds, (ii) adsorption of SiH_3 by reaction with surface Si dangling bonds, (iii) breaking of Si-Si dimer bonds by SiH_3 and subsequent formation of dihydrides on the surface, and (iv) formation of disilane on the surface by reaction between neighboring SiH_3 species attached to the surface. Further trajectory generation and visualization over the entire temperature range of interest will be carried out during the next fiscal year.
- Detailed analysis of the reactions identified in the MD trajectories as described in the previous part. Some of the key reactions that have been identified during the initial stage of deposition have been analyzed. The remaining reactions of the initial stage and the key

thermally activated processes of the later stages of deposition will be analyzed during the next fiscal year.

- Reaction rate calculations for the reactions analyzed above (as described in the previous part for the isolated reactions) and comparison with the rates that correspond to the isolated reactions. This work will be completed during the next fiscal year: it will provide the necessary input and determine the range of validity of certain simplifying assumptions that can be used in the KMC simulations.

MD simulation studies with additional deposition precursors to SiH_3 will be initiated during the next fiscal year. The relative rates of impingement on the surface of the different species in the MD simulations will be determined by experimental species flux measurements and in conjunction with plasma-phase fluid models. In addition, a systematic analysis of the substrate temperature effects and the detailed structural and chemical characterization of the deposited films for all of the different sets of deposition precursors will be performed during the next fiscal year. Finally, tasks that are anticipated to be initiated during the next fiscal year include the implementation of KMC deposition simulations based on the available reaction rate database, as well as the empirical interatomic description of the Si:H:Cl system, i.e., the additional incorporation of Cl into the radical-surface chemistry for modeling plasma deposition from dichlorosilane (SiH_2Cl_2) containing discharges.

RESEARCH IN AERODYNAMIC SHAPE OPTIMIZATION

James Reuther, Ph.D (RIACS)

Since the inception of CFD, researchers have sought not only accurate aerodynamic prediction methods for given configurations, but also design methods capable of creating new optimum configurations. Yet, while flow analysis can now be carried out over quite complex configurations using the Navier-Stokes equations with a high degree of confidence, direct CFD based design is still a daunting challenge for complex three-dimensional problems. This is especially true in problems where viscous effects play a dominant role. The main effort of this research is the introduction of new technology to overcome the difficulties present in traditional aerodynamic optimization methods. The CFD-based aerodynamic design methods of the past can be grouped into two basic categories: inverse methods, and numerical optimization methods.

Inverse methods derive their name from the fact that they invert the goal of the flow analysis algorithm. Instead of obtaining the surface distribution of an aerodynamic quantity, such as pressure, for a given shape, they calculate the shape for a given surface distribution of an aerodynamic quantity. Most of these methods are based on potential flow techniques, and few of them have been extended to three-dimensions. The common trait of all inverse methods is their computational efficiency. Typically, transonic inverse methods require the equivalent of 2-10 complete flow solutions in order to render a complete design. Since obtaining a few solutions for simple two-dimensional and three-dimensional designs can be done in at most a few hours on modern computer systems, the computational cost of most inverse methods is considered to be minimal. Unfortunately, they suffer from many limitations and difficulties the most glaring of which is that the objective is built directly into the design process and thus cannot be changed to an arbitrary or more appropriate objective function.

A traditional alternative, which avoids some of the difficulties of inverse methods while incurring a heavy computational expense, is the use of numerical optimization methods. The

essence of these methods is straightforward: a numerical optimization procedure is coupled directly to an existing CFD analysis algorithm. The numerical optimization procedure attempts to extremize a chosen aerodynamic measure of merit, which is evaluated by the chosen CFD code. Most of these optimization procedures require gradient information in addition to evaluations of the objective function. Here, the gradient refers to changes in the objective function with respect to changes in the design variables. The simplest method of obtaining gradient information is by finite differences. In this technique, the gradient components are estimated by independently perturbing each design variable with a finite step, calculating the corresponding value of the objective function using CFD analysis, and forming the ratio of the differences. These methods are very versatile, allowing any reasonable aerodynamic quantity to be used as the objective function. They can be used to mimic an inverse method by minimizing the difference between target and actual pressure distributions, or may instead be used to maximize other aerodynamic quantities of merit such as L/D . Unfortunately, these finite difference numerical optimization methods, unlike the inverse methods, are computationally expensive because of the large number of flow solutions needed to determine the gradient information for a useful number of design variables. Tens of thousands of flow analyses would be required for a complete three-dimensional design.

In this research, a new method is developed that avoids the limitations and difficulties of traditional inverse methods while retaining their inherent computational efficiency. The method dramatically reduces the cost of aerodynamic optimization by replacing the expensive finite-difference method of calculating the required gradients with an adjoint variable formulation. After deriving the differential form of the adjoint equations and posing the correct boundary conditions based on the objective function, the resulting system is discretized and solved on the same mesh as that used for the flow solution. A significant economization is thus achieved by applying many of the subroutines used for the flow solution to the solution of the adjoint equations. The resulting design process requires only one flow calculation and one adjoint calculation per gradient evaluation, as opposed to the hundreds required for a finite-difference gradient involving hundreds of design variables. In practice the computational cost of the new method is two orders of magnitude less than that of a conventional approach. Considerable effort has been focused in the last two years to develop control theory-based aerodynamic shape optimization methods. This effort has been conducted by a team of researchers from around the nation whose major contributors include Prof. Antony Jameson of Stanford University, Prof. Luigi Martinelli of Princeton University, Prof. Juan J. Alonso of Stanford University, Dr. James Farmer and myself. Many of the core subroutines upon which the research has been formulated are the intellectual property of Intelligent Aerodynamics International. The work that has taken place in the last three years can be broken down into three specific areas.

- 1) Two-dimensional and three-dimensional proof-of-concept studies.
- 2) The development and demonstration of a three-dimensional research tool for complex configurations.
- 3) NASA and industrial evaluation and feedback.

During the first year, work was primarily focused in area (1) and to a lesser extent areas (2) and (3). At the beginning of this program at RIACS, methods were in place, which showed that control theory could be used in conjunction with numerical optimization and computational fluid dynamics to create efficient design tools for flows governed by the potential flow equation (AIAA Paper 94-0499).

During the course of the first year of the program the development of adjoint methods was extended to treat the Euler equations. In our paper at the Multi-Disciplinary Optimization conference during summer 1994 (AIAA paper 94-4272, also RIACS report 94.18), results were shown demonstrating that control theory based on the Euler equations could be used to design airfoils that operate under transonic conditions. Various objective functions were demonstrated

showing the versatility of the new method. In the work presented at VKI, the first examples of three-dimensional wing design using control theory were presented. Finally, in a paper presented at the January 1995 Aerospace Sciences Meeting (AIAA paper 95-0123, also RIACS report 95.01) results for the design of wing and wing-body configurations over general meshes were shown.

One of the dramatic successes in the first year involved the participation of the Beechcraft Aircraft Division of Raytheon, Inc. Raytheon entered into a cooperative agreement with NASA Ames Research Center to explore the usefulness of the adjoint-based design optimization methods. Between March and May of 1995, a team of scientists from Raytheon and Ames were able to combine their talents by employing a preliminary version of the three-dimensional design code, described in RIACS report 95.01, to develop a new wing for the Premier I business jet configuration. This one-month design of a new transonic wing contrasts with the usual development time of more than a year for traditional methods. Raytheon has since conducted wind tunnel tests confirming that the new wing design realizes its predicted performance and launched the design for production. They subsequently took 51 orders for the new airplane on the day the design was announced. Furthermore, Raytheon has been so impressed by the capability of adjoint-based design methods that they are now incorporating them into their own aircraft design environments. A paper authored by both NASA and Raytheon personnel that presents the basic design strategy and its outcome was presented at the Aerospace Sciences Meeting, January 1996 (AIAA paper 96-0554, also RIACS 96.03).

Another group that has taken a keen interest in our research from its first year is the NASA High Speed Research Program (HSR) group. In their effort to create economically viable supersonic transport configurations for the next century they are investigating the use of aerodynamic shape optimization to improve aerodynamic performance for the High Speed Civil Transport. Both the traditional and the adjoint-based design methods studied by our group at Ames have been tested by the HSR community. A paper that gives an example of the capabilities of this emerging technology for supersonic design was presented at the American Society of Mechanical Engineers annual winter meeting in November 1995 (also RIACS 95.14).

At the beginning of the second year of this research, experience with both Raytheon and the HSR group highlighted the need to develop an enhanced implementation of the aerodynamic shape optimization method, which would allow the treatment of more complex geometries. Until that time, only a single-grid-block design method had been developed, capable of handling wing/body configurations but leaving engine nacelle effects to be modeled by approximations at best. However, the real-life problems presented by industry required the design method to handle complete aircraft configurations, which in turn mandated either an extension to a multiblock grid topology or a switch to unstructured meshes. The former path was chosen because of its relatively straightforward implementation and the natural avenue, which multiple blocks provide towards a parallelization of the process.

The first paper demonstrating the new multiblock capability was presented at the 34th Aerospace Sciences Meeting (AIAA paper 96-0094, RIACS report 96.02). Following this paper, the focus quickly turned to a parallel form of the multiblock software. This was essential because the added complexity of complete aircraft configurations required a significant increase in the harnessed computer power. A paper presented at the Multi-Disciplinary Optimization Conference in September 1996 highlighted this parallel multiblock capability.

The second year was also characterized by a significant effort to enhance the software for the HSR work. Since details of this work cannot be presented in view of its sensitive nature, it must suffice to state that both the single-block and multiblock codes were modified so that HSR-specific design problems could be treated robustly and efficiently. One major activity was the incorporation of a constrained optimization capability as opposed to the use of an

unconstrained algorithm. This was necessitated by the hundreds of geometric constraints imposed on the HSRP configurations (such as on wing spar thickness, fuel volume, and cabin dimensions). The year culminated in the successful application of the HSRP-specific versions of the software to an industry-established test-bed configuration. Independent methods were also applied to the same problem by Boeing and McDonnell Douglas teams. This constrained optimization exercise showed the software used at NASA Ames to be effective and favorable. Further efforts during this last year have focused on enhancing the multiblock design capability to treat even more realistic problems. The first step in this path was the inclusion of constraints and the treatment of multiple design points. A paper presented at the 35th Aerospace Sciences Meeting (AIAA paper 97-0103, RIACS report 97.02) highlighted these capabilities. The next step was the inclusion of a viscous design capability through the extension of the underlying flow solver from the Euler equations to the Navier-Stokes equations. An important point is that the solution cost in terms of computer time to solve the Navier-Stokes equations as opposed to the Euler equations is a factor of about 5. Further, since multiple flow solutions are required to solve a design problem, the use of parallel computing, first introduced in the second year of this research, has become an essential capability of the software. Navier-Stokes-based design problems typically require on the order of 32 SGI Origin 2000 CPUs for roughly 24 hours. This level of computer resources corresponds to computer turn-around times that would be unacceptable on the fastest serial CPUs that are available today. Considerable time was thus invested to ensure that the software ran robustly in parallel and on various platforms. To date, the parallel multiblock code has been ported to the IBM SP2, the CRAY J90 and C90, and SGI Origin 2000, and a cluster of HP workstations. The need for parallel efficiency on these widely varying architectures required careful management and tuning of the interprocessor communication costs. Issues included load balancing, bandwidth minimization, latency reduction, and scalability with respect to total mesh size. A paper presented at the AIAA 13th Biannual CFD Conference (AIAA paper 97-1893 RIACS report 97.05) discussed both the extension to the Navier-Stokes equations and the details related to improving the parallel performance.

In the last and final year of this program, work consisted of ensuring that the capabilities developed within the program were available for use in NASA programs. Work with the HSRP program included the inclusion of a more sophisticated geometry treatment package. This new package which has been integrated into the multiblock version of the software will allow the simultaneous design of a wing and closely coupled nacelles. With an enhanced multipoint capability, new engine inlet and exit boundary conditions, and a new treatment of aerodynamic constraints, the improved software is now prepared for use on the difficult design problems foreseen in the coming year of the HSR program. The collaborative work with Steve Edwards, a Masters level student from UC Davis, has been concluded with the completion of his thesis and a new the capability of performing efficient shape optimization of airfoils at transonic speeds including viscous effects.

The research involved using a coupled Euler + boundary layer formulation combined with a similarly coupled viscous adjoint solver. Other work investigating the viscous design problem which involved the inclusion of other turbulence models into the existing capability has progressed but will have to be continued outside the scope of this RIACS project. A final area where new research was explored in the last year consists of development of an aeroelastic shape design capability. As it turns out, much of the work that was necessary in order to arrive at an aerodynamic design capability for complex configurations can also be used for aeroelastic analysis or design. Thus work was begun to develop a general aeroelastic analysis/design capability which can accommodate other disciplines as necessary. The general framework has been completed, however the detailed implementation will have to be completed outside the present RIACS contract. Three Journal articles based on the work over the last three years were also completed and accepted for publication. Two will be published in the AIAA Journal of Aircraft while a third will be published in the Journal of Computers and Fluids.

While this three-year project has only served to begin the process of developing advanced aerodynamic shape optimization capabilities it has been a rewarding effort. The methods that have been the topic of study have seen wide spread use both inside and outside of NASA. The group of collaborative researchers that have together advances the theoretical and implementational accomplishments highlighted in the above statement have received substantial acclaim for pushing a new technology that is now recognized as an independent sub-discipline within the wider CFD field. The developments achieved thus far have demonstrated beyond a doubt the great value of adjoint-based aerodynamic design. It is hoped that with all of these advances, the greater aeronautical science community will in the future adopt these new ideas into their production design environments. Certainly if the work in conjunction with Raytheon is any indication, this is already taking place.

X-Zm-Content-Name: abstract-jair1.tex

Content-Description: Text

Content-Type: text/plain; name="abstract-jair1.tex"; charset=us-ascii

The following is the first of a two-paper series that describes a shape optimization method for aerodynamic design problems involving complex aircraft configurations, multiple design points, and subject to geometric constraints. The aerodynamic performance is evaluated using a set of high fidelity governing equations discretized on body-conforming multiblock meshes. The design process is greatly accelerated through the use of an adjoint method for the calculation of sensitivity information and by means of a parallel implementation for distributed memory computers.

This first paper of the series focuses on the details of the development and implementation of the complete design algorithm. A general adjoint formulation of the design optimization problem is presented along with a detailed derivation of the adjoint system for the Euler equations. Both the multiblock approaches for the flow and adjoint solution algorithm and the mesh perturbation procedure are described. The extension of this approach to treat multipoint and constrained shape optimization of complete aircraft configurations is also discussed. Finally, the details of the parallel implementation are examined. The second paper of the series explores specific design examples and assesses the performance of the design algorithm.

X-Zm-Content-Name: abstract-jair2.tex

Content-Description: Text

Content-Type: text/plain; name="abstract-jair2.tex"; charset=us-ascii

This is the second of a two-paper series that describes the derivation, development, and practical application of a shape optimization method for aerodynamic design problems involving complex aircraft configurations including multiple design points and subject to geometric constraints. The aerodynamic performance is evaluated using the Euler equations discretized on body-conforming multiblock meshes. The design process is greatly accelerated by both the use of control theory for the calculation of sensitivity information and the implementation of the method on distributed memory parallel computers.

While the first paper of the series placed emphasis on the foundations of the control theory approach and presented a general framework for the design of complex configurations, this second paper focuses on both the performance of the overall design method and its practical application to realistic design problems. More specifically, the accuracy of the adjoint-based sensitivity information is examined and compared with the sensitivities provided by the finite difference method. The parallel scalability of the design method is analyzed and demonstrated. Aerodynamic performance improvements are obtained for example wing design problems using both multipoint drag minimization and single point inverse techniques. Design examples

include both transonic and supersonic cases and involve wing aerodynamic design in the presense of complex aircraft configurations with all resulting configurations satisfying geometric constraints.

SIMULATION OF REAL-WORLD PROBLEMS

Andrew Sohn, Ph.D (RIACS)

Simulations of real-world problems are a challenging task. Realistically sized problems often take a tremendous amount of computational power. The current approach to solve real-world computational science problems is the use of large-scale distributed-memory multiprocessors. However, this approach requires much fundamental research in computer science because the conventional methods that work for small scale machines do not work on large-scale machines. This difficulty in using large-scale machines necessitates much research into parallel algorithms, architectures, compilers, and its programming environment. The Information Power Grid (IPG) project currently being undertaken at NASA Ames aims at performing research in basic issues towards solving real-world problems on large-scale machines in a consolidated and effective way.

RIACS research interests are in parallel computing, communication studies of distributed-memory multiprocessors, and multithreaded architectures and its compilers. The main goal of research under RIACS travel support is to investigate the feasibility of scalable computing towards the IPG project. To achieve the objective, we have been designing and implementing a set of representative computer science problems on a set of representative large-scale distributed-memory multiprocessors. The target problems include sorting, graph partitioning, and search problems that are considered fundamental problems in computer science regardless of their applications. The target machines include the Cray/SGI T3E, IBM SP-2, SGI Origin, and ETL EM-X. All these machines have more than 64 processors.

As of now, RIACS has been able to design and implement a new radix sorting algorithm, called balanced radix sort. Experimental results on large-scale parallel machines indicate that the algorithm is currently the fastest among parallel sorting algorithms, about 1.3 times to twice faster than sample sorting, which used to be the fastest sorting algorithm. RIACS has also designed and implemented a fast graph Partitioning algorithm together with Horst Simon of NERSC of Lawrence Berkeley Laboratory, called S-HARP. Experimental results also indicate that the algorithm is the fastest among parallel partitioning algorithms, about 3 to 10 times faster than the fastest one.

RIACS has just begun this practical approach to realize part of the IPG project. There are a lot more issues to be investigated than addressed. For sorting, there are two more sorting algorithms that we intend to design, develop and implement on the parallel machines. These algorithms include bitonic sorting which is the simplest in communication between processors and sample sorting which used to be the fastest sorting algorithm. For graph partitioning, there are numerous issues that need to be addressed. We intend to perform as much as we can while at RIACS under travel support. The approach RIACS has been undertaking is perhaps one of the most practical yet feasible approaches to realize the IPG project because the results will allow us to compare the relative performance of large-scale distributed-memory machines.

III. TECHNICAL REPORTS

98.01 PLUM: PARALLEL LOAD BALANCING FOR UNSTRUCTURED ADAPTIVE MESHES, *Leonid Oliker, January 09, 1998 (121 pages) (Defended at the University of Colorado, November 13, 1997.)*

Dynamic mesh adaptation on unstructured grids is a powerful tool for computing large-scale problems that require grid modifications to efficiently resolve solution features. Unfortunately, an efficient parallel implementation is difficult to achieve, primarily due to the load imbalance created by the dynamically changing nonuniform grid. To address this problem, we have developed PLUM, an automatic portable framework for performing adaptive large-scale numerical computations in a message-passing environment.

First, we present an efficient parallel implementation of a tetrahedral mesh adaptation scheme. Extremely promising parallel performance is achieved for various refinement and coarsening strategies on a realistic-sized domain. Next we describe PLUM, a novel method for dynamically balancing the processor workloads in adaptive grid computations. This research includes interfacing the parallel mesh adaptation procedure based on actual flow solutions to a data remapping module, and incorporating an efficient parallel mesh repartitioner. A significant runtime improvement is achieved by observing that data movement for a refinement step should be performed after the edge-marking phase but before the actual subdivision. We also present optimal and heuristic remapping cost metrics that can accurately predict the total overhead for data redistribution.

Several experiments are performed to verify the effectiveness of PLUM on sequences of dynamically adapted unstructured grids. Portability is demonstrated by presenting results on the two vastly different architectures of the SP2 and the Origin2000. Additionally, we evaluate the performance of five state-of-the-art-partitioning algorithms that can be used within PLUM. It is shown that for certain classes of unsteady adaptation, globally repartitioning the computational mesh produces higher quality results than diffusive repartitioning schemes. We also demonstrate that a coarse starting mesh produces high quality load balancing, at a fraction of the cost required for a fine initial mesh. Results indicate that our parallel load balancing strategy will remain viable on large numbers of processors.

98.02: PERFORMANCE ANALYSIS AND PORTABILITY OF THE PLUM LOAD BALANCING SYSTEM, *Leonid Oliker, Rupak Biswas (NASA Ames), and Harold N. Gabow (University of Colorado), September 1998 (10 pages)*

The ability to dynamically adapt an unstructured mesh is a powerful tool for solving computational problems with evolving physical features; however, an efficient parallel implementation is rather difficult. To address this problem, we have developed PLUM, an automatic portable framework for performing adaptive numerical computations in a message-passing environment. PLUM requires that all data be globally redistributed after each mesh adaptation to achieve load balance. We present an algorithm for minimizing this remapping overhead by guaranteeing an optimal processor reassignment. We also show that the data redistribution cost can be significantly reduced by applying our heuristic processor reassignment algorithm to the default mapping of the parallel partitioner. Portability is examined by comparing performance on a SP2, an Origin2000, and a T3E. Results show that PLUM can be successfully ported to different platforms without any code modifications.

98.03 : AUTOMATING MISSION SCHEDULING FOR SPACE-BASED OBSERVATORIES *Nicola Muscettola (Recom), Barney Pell (RIACS), Othar Hansson (Heuristicrats Research, Inc), Sunil Mohan (Recom), March 16, 1998 (14 pages)*

(Appeared in G.W. Henry and J.A. Eaton, eds, "Robotic Telescopes: Current Capabilities, Present Developments, and Future Prospects for Automated Astronomy", Provo, UT: Astronomical Society of the Pacific, ASP Conference Series number 79, 1995.)

In this paper we describe the use of our planning and scheduling framework, HSTS, to reduce the complexity of science mission planning. This work is part of an overall project to enable a small team of scientists to control the operations of a spacecraft. The present process is highly labor intensive. Users (scientists and operators) rely on a non-codified understanding of the different spacecraft subsystems and of their operating constraints. They use a variety of software tools to support their decision making process. This paper considers the types of decision making that need to be supported/automated, the nature of the domain constraints and the capabilities needed to address them successfully, and the nature of external software systems with which the core planning/scheduling engine needs to interact. HSTS has been applied to science scheduling for EUVE and Cassini and is being adapted to support autonomous spacecraft operations in the New Millennium initiative.

98.04 : A HYBRID PROCEDURAL/DEDUCTIVE EXECUTIVE FOR AUTONOMOUS SPACECRAFT *Barney Pell (RIACS), Ed Gamble (JPL), Erann Gat (JPL), Ron Keesing (Caelum Research), Jim Kurien, (NASA Ames) Bill Millar (Caelum Research), P. Pandurang Nayak (RIACS), Christain Plaunt, (Caelum Research) and Brian Williams (NASA Ames).*

(To appear in the Proceedings of the Second International Conference on Autonomous Agents, ACM Press, 1998.)

The New Millennium Remote Agent (NMRA) will be the first AI system to control an actual spacecraft. The spacecraft domain places a strong premium on autonomy and requires dynamic recoveries and robust concurrent execution, all in the presence of tight real-time deadlines, changing goals, scarce resource constraints, and a wide variety of possible failures. To achieve this level of execution robustness, we have integrated a procedural executive based on generic procedures with a deductive model-based executive. A procedural executive provides sophisticated control constructs such as loops, parallel activity, locks, and synchronization which are used for robust schedule execution, hierarchical task decomposition, and routine configuration management. A deductive executive provides algorithms for sophisticated state inference and optimal failure recovery planning. The integrated executive enables designers to code knowledge via a combination of procedures and declarative models, yielding a rich modeling capability suitable to the challenges of real spacecraft control. The interface between the two executives ensures both that recovery sequences are smoothly merged into high-level schedule execution and that a high degree of reactivity is retained to effectively handle additional failures during recovery.

98.05 : ISSUES IN TEMPORAL REASONING FOR AUTONOMOUS CONTROL SYSTEMS *Nicola Muscettola (Recom), Paul Morris (Caelum Research), Barney Pell (RIACS), and Ben Smith (JPL)*

(To appear in the Proceedings of the Second International Conference on Autonomous Agents, ACM Press, 1998)

Deep Space One will be the first spacecraft to be controlled by an autonomous agent potentially capable of carrying out a complete mission with minimal commanding from Earth. The New Millennium Remote Agent (NMRA) includes a planner-scheduler that produces plans, and an executive that carries them out. In this paper we discuss several issues arising at the interface between planning and execution. Issues include execution latency, plan dispatchability, and the distinction between controllable and uncontrollable events. Temporal information in the plan is

represented within the general framework of Simple Temporal Constraint networks, as introduced by Dechter, Meiri, and Pearl. However, the execution requirements have a substantial impact on the topology of the links and the propagation through the network.

98.06 : DESIGN OF THE REMOTE AGENT EXPERIMENT FOR SPACECRAFT AUTONOMY

P. Pandurang Nayak (RIACS), Barney Pell (RIACS), Douglas E. Bernard (JPL), Gregory A. Dorais (Caelum Research), Chuck Fry (Caelum Research), Edward Gamble Jr. (JPL), Bob Kanefsky (Caelum Research), James Kurien (Caelum Research), William Millar (Caelum Research), Nicola Muscettola (Recom Technologies), Kanna Rajan (Caelum Research), Nicolas Rouquette (JPL), Benjamin Smith (JPL), Brian C. Williams (NASA)

This paper describes the Remote Agent flight experiment for spacecraft commanding and control. In the Remote Agent approach, the operational rules and constraints are encoded in the flight software. The software may be considered to be an autonomous "remote agent" of the space operators in the sense that the operators rely on the agent to achieve particular goals.

The experiment will be executed during the flight of NASA's Deep Space One technology validation mission. During the experiment, the spacecraft will not be given the usual detailed sequence of commands to execute. Instead, the spacecraft will be given a list of goals to achieve during the experiment. In flight, the Remote Agent flight software will generate a plan to accomplish the goals and then execute the plan in a robust manner while keeping track of how well the plan is being accomplished. During plan execution, the Remote Agent stays on the lookout for any hardware faults that might require recovery actions or replanning.

In addition to describing the design of the remote agent, this paper discusses technology-insertion challenges and the approach used in the Remote Agent approach to address the challenges.

The experiment integrates several spacecraft autonomy technologies developed at NASA Ames and the Jet Propulsion Laboratory: on-board planning, a robust multi-threaded executive, and model-based failure diagnosis and recovery.

98.07 : THE REMOTE AGENT EXECUTIVE: CAPABILITIES TO SUPPORT INTEGRATED ROBOTIC AGENTS *Barney Pell (RIACS), Gregory A. Dorais (Caelum Research), Christian Plaunt (Caelum Research) and Richard Washington (Caelum Research)*

(In Schultz, A., and Kortenkamp, D., eds., *Proceeds of the AAAI Spring Symposium. on Integrated Robotic Architectures*. Palo Alto, CA: AAAI Press, 1998)

The Remote Agent (RA) integrates a broad spectrum of robotic capabilities, including planning, scheduling, execution, monitoring, failure detection, diagnosis, and recovery. The RA Executive (EXEC) can be viewed as the core of the agent. The capabilities built into EXEC have been useful in moving from a set of separate components to a fully integrated robotic agent. Since the abstract capabilities enable software developers to think about the robot at a higher level, these capabilities also support the reuse of knowledge and code across multiple robot applications. These capabilities include a high-level procedural action-description language, a multi-threaded operating system, services for resource management and configuration management, support for closed-loop plan execution and replanning, routines for executing routine low-level command sequences, a framework for specifying fault responses, and a way to implement a top-level loop including behavior during initialization, safe modes, and responses to plan failures.

We believe that these capabilities are required for most autonomous agents, and that executives and execution capabilities will become more essential as we attempt to develop autonomous

agents of increasing capability. To this end, we are modularizing the executive by providing each capability as a separate module so that the individual capabilities can be standardized and shared across different agent architectures.

98.08 : MISSION OPERATIONS WITH AN AUTONOMOUS AGENT

B. Pell (RIACS) S. Sawyer (Lockheed), N. Muscettola (Recom), B. Smith (JPL), and D.E. Bernard (JPL).

(In Proceedings of the IEEE Aerospace Conference, March, 1998.)

The Remote Agent (RA) is an Artificial Intelligence (AI) system which automates some of the tasks normally reserved for human mission operators and performs these tasks autonomously on-board the spacecraft. These tasks include activity generation, sequencing, spacecraft analysis, and failure recovery. The RA will be demonstrated as a flight experiment on Deep Space One (DS1), the first deep space mission of the NASA's New Millennium Program (NMP). As we moved from prototyping into actual flight code development and teamed with ground operators, we made several major extensions to the RA architecture to address the broader operational context in which RA would be used. These extensions support ground operators and the RA sharing a long-range mission profile with facilities for asynchronous ground updates; support ground operators monitoring and commanding the spacecraft at multiple levels of detail simultaneously; and enable ground operators to provide additional knowledge to the RA, such as parameter updates, model updates, and diagnostic information, without interfering with the activities of the RA or leaving the system in an inconsistent state.

The resulting architecture supports incremental autonomy, in which a basic agent can be delivered early and then used in an increasingly autonomous manner over the lifetime of the mission. It also supports variable autonomy, as it enables ground operators to benefit from autonomy when they want it, but does not inhibit them from obtaining a detailed understanding and exercising tighter control when necessary. These issues are critical to the successful development and operation of autonomous spacecraft.

98.09 : REMOTE AGENT: TO BOLDLY GO WHERE NO AI SYSTEM HAS GONE BEFORE *N. Muscettola (Recom), P.P. Nayak (RIACS) , B. Pell (RIACS), and B.C. Williams (NASA Ames).*

(Submitted to Artificial Intelligence, 1998)

Renewed motives for space exploration have inspired NASA to work toward the goal of establishing a virtual presence in space, through heterogeneous fleets of robotic explorers. Information technology, and Artificial Intelligence in particular, will play a central role in this endeavor by endowing these explorers with a form of computational intelligence that we call "Remote agents". In this paper we describe the Remote Agent, a specific autonomous agent architecture based on the principles of model-based programming, on-board deduction and search, and goal-directed closed-loop commanding, that takes a significant step toward enabling this future. This architecture addresses the unique characteristics of the spacecraft domain that require highly reliable autonomous operations over long periods of time with tight deadlines, resource constraints, and concurrent activity among tightly coupled subsystems. The Remote Agent integrates constraint-based temporal planning and scheduling, robust multi-threaded execution, and model-based mode identification and reconfiguration. The demonstration of the integrated system as an on-board controller for Deep Space One, NASA's first New Millennium mission, is scheduled for a period of a week in late 1998. The development of the Remote Agent also provided the opportunity to reassess some of AI's conventional wisdom about the challenges of implementing embedded systems, tractable reasoning, and knowledge

representation. We discuss these issues, and our often contrary experiences, throughout the paper.

98.10 : PERFORMANCE OF LOW DISSIPATIVE HIGH ORDER SHOCK-CAPTURING SCHEMES FOR SHOCK-TURBULENCE INTERACTIONS, *N.D. Sandham and H.C. Yee*, April 1998 (44 pages)

(Invited paper for the 6th ICFD Conference on Numerical Methods for Fluid Dynamics University of Oxford, Oxford, England)

Accurate and efficient direct simulation of turbulence in the presence of shock waves represents a significant challenge for numerical methods. The objective of this paper is to evaluate the performance of high order compact and non-compact central spatial differencing employing total variation diminishing (TVD) shock-capturing schemes as characteristic based filters for two model problems combining shock wave and shear layer phenomena. A vortex pairing model evaluates the ability of the schemes to cope with shear layer instability and eddy shock waves, while a shock wave impingement on a spatially-evolving mixing layer model studies the accuracy of computation of vortices passing through a sequence of shock and expansion waves. A drastic increase in accuracy is observed if a suitable artificial compression formulation is applied to the TVD schemes. With this modification to the filter step the fourth order non-compact scheme shows improved results in comparison to second order methods, while retaining the good shock resolution of the basic TVD scheme. For this characteristic based filter approach, however, the benefits of compact schemes or schemes with higher than fourth order are not sufficient to justify the higher complexity near the boundary and/or computational cost.

98.11 : LOW DISSIPATIVE HIGH ORDER SHOCK-CAPTURING METHODS USING CHARACTERISTIC-BASED FILTERS, *H.C. Yee, N.D. Sandham and M.J. Djomehri* May 1998 (58 pages)

(Submitted to Journal of Computational Physics)

An approach which closely maintains the non-dissipative nature of classical fourth or higher-order spatial differencing away from shock waves and high gradient regions while being capable of accurately capturing discontinuities, high gradient and fine scale turbulent structures in a stable and efficient manner is described. The approach is a generalization of the method of Gustafsson and Olsson and the artificial compression method (ACM) of Harten. Spatially non-dissipative fourth-order or higher compact and non-compact spatial differencings are used as the base schemes. Instead of applying a scalar filter as in Gustafsson and Olsson, an ACM like term is used to signal the appropriate amount of second or third-order TVD or ENO types of characteristic based numerical dissipation. This term acts as a characteristic filter to minimize numerical dissipation for the overall scheme. For time-accurate computations, time discretizations with low dissipation are used. Numerical experiments on 2-D vertical flows, vortex-shock interactions and compressible spatially and temporally evolving mixing layers showed that the proposed schemes have the desired property with only a 10% increase in operations count over standard second-order TVD schemes. Aside from the ability to accurately capture shock-turbulence interaction flows, this approach is also capable of accurately preserving vortex convection. Higher accuracy is achieved with fewer grid points when compared to that of standard second-order TVD or ENO schemes. To demonstrate the applicability of these schemes in sustaining turbulence where shock waves are absent, a simulation of 3-D compressible turbulent channel flow in a small domain is conducted.

98.12 : PLUM: PARALLEL LOAD BALANCING FOR UNSTRUCTURED ADAPTIVE MESHES
Leonid Oliker (RIACS), Rupak Biswas (NASA Ames), June 1998 (27 pages)

(To Appear in the Journal of Parallel and Distributed Computing (JPDC), 1998.)

Mesh adaption is a powerful tool for efficient unstructured-grid computations but causes load imbalance among processors on a parallel machine. We present a novel method called PLUM to dynamically balance the processor workloads with a global view. This paper describes the implementation and integration of all major components within our dynamic load balancing strategy for adaptive grid calculations. Mesh adaption, repartitioning, processor assignment, and remapping are critical components of the framework that must be accomplished rapidly and efficiently so as not to cause a significant overhead to the numerical simulation. A data redistribution model is also presented that predicts the remapping cost on the SP2. This model is required to determine whether the gain from a balanced workload distribution offsets the cost of data movement. Results presented in this paper demonstrate that PLUM is an effective dynamic load balancing strategy which remains viable on a large number of processors.

98.13 : CONSTRAINED MULTI-POINT AERODYNAMIC SHAPE OPTIMIZATION USING AN ADJOINT FORMULATION AND PARALLEL COMPUTERS, PART 1

James Reuther (RIACS), Anthony Jameson, Juan Jose Alonso, Mark J. Rimlinger (Raytheon) and David Saunders (Raytheon), July 1998 (37 pages)

(Accepted for publication AIAA Journal of Aircraft)

The following is the first of a two-paper series that describes a shape optimization method for aerodynamic design problems involving complex aircraft configurations, multiple design points, and subject to geometric constraints. The aerodynamic performance is evaluated using a set of high fidelity governing equations discretized on body-conforming multiblock meshes. The design process is greatly accelerated through the use of an adjoint method for the calculation of sensitivity information and by means of a parallel implementation for distributed memory computers.

This first paper of the series focuses on the details of the development and implementation of the complete design algorithm. A general adjoint formulation of the design optimization problem is presented along with a detailed derivation of the adjoint system for the Euler equations. Both the multiblock approach for the flow and adjoint solution algorithm and the mesh perturbation procedure are described. The extension of this approach to treat multipoint and constrained shape optimization of complete aircraft configurations is also discussed. Finally, the details of the parallel implementation are examined. The second paper of the series explores specific design examples and assesses the performance of the design algorithm.

98.14: CONSTRAINED MULTIPOINT AERODYNAMIC SHAPE OPTIMIZATION USING AN ADJOINT FORMULATION AND PARALLEL COMPUTERS: PART II. *James Reuther (RIACS), Anthony Jameson, Juan Jose Alonso (Stanford University), Mark J. Rimlinger (Sterling Software) and David Saunders (Sterling Software), July 1998 (40 pages)*

This is the second of a two-paper series that describes the derivation, development, and practical application of a shape optimization method for aerodynamic design problems involving complex aircraft configurations including multiple design points and subject to geometric constraints. The aerodynamic performance is evaluated using the Euler equations discretized on body-conforming multiblock meshes. The design process is greatly accelerated by both the use of control theory for the calculation of sensitivity information and the implementation of the method on distributed memory parallel computers.

While the first paper of the series placed emphasis on the foundations of the control theory approach and presented a general framework for the design of complex configurations, this second paper focuses on both the performance of the overall design method and its practical application to realistic design problems. More specifically, the accuracy of the adjoint-based sensitivity information is examined and compared with the sensitivities provided by the finite difference method. The parallel scalability of the design method is analyzed and demonstrated. Aerodynamic performance improvements are obtained for example wing design problems using both multipoint drag minimization and single point inverse techniques. Design examples include both transonic and supersonic cases and involve wing aerodynamic design in the presence of complex aircraft configurations with all resulting configurations satisfying geometric constraints.

98.15: A PERFORMANCE STUDY OF DIFFUSIVE VS. REMAPPED LOAD-BALANCING SCHEMES, *K.Schloegel (Dept. of Science & Engineering) G.Karypis (Dept. of Science & Engineering), V.Kumar (Dept. of Science & Engineering), R.Biswas (NASA Ames) and L.Oliker (RIACS), August 1998 (8 pages)*

(To appear in the proceedings of the 11th International Conference on Parallel and Distributed Computing Systems 1998.)

For a large class of irregular grid applications, the computational structure of the problem changes in an incremental fashion from one phase of the computation to another. Eventually, as the graph evolves, it becomes necessary to correct the partition in accordance with the structural changes in the computation. Partitioning the graph from scratch and then intelligently remapping the resulting partition will accomplish this task. Two different types of schemes to accomplish this task have been developed recently. In one scheme, the graph is partitioned from scratch and then the resulting partition is remapped intelligently to the original partition. The second type of scheme use a multilevel diffusion repartitioner. In this paper, we conduct a comparison study on repartitioning via intelligent remapping versus repartitioning by diffusion.

We show that multilevel diffusion algorithms generally produce significantly lower data migration overhead for adaptive graphs in which low magnitude localized or non-localized imbalances occur and for graphs in which high magnitude imbalances occurs globally throughout the domains than partitioning from scratch and remapping the resulting partition.

We show that for the class of problems in which high magnitude imbalances occur in localized areas of the graph, partitioning from scratch and remapping the resulting partition will result in very low edge-cuts and data migration overheads which are similar to those obtained by diffusive schemes.

Finally, we show that the run times of the various schemes are all similar.

IV. PUBLICATIONS

P. Cheeseman, D.G. Balwin, and W.J. Emery, "Higher Resolution Earth Surface Features from Repeat Moderate Resolution Satellite Imagery". IEEE Trans. On Geo. And Rem. Sens., Vol. 36, No. 1, Jan. 1998

B. Pell, E. Gamble, E. Gat, R. Keesing, J. Kurien, W. Millar, P. Nayak, C. Plaunt, B. Williams. "A Hybrid Procedural/Deductive Executive For Autonomous Spacecraft." To appear in Journal of Autonomous Agents and Multi-Agent Systems 2(1), 1999.

P. Nayak, B. Pell, D. Bernard, G. Dorais, C. Fry, E. Gamble, R. Kanefsky, J. Kurien, W. Millar, N. Muscettola, K. Rajan, N. Rouquette, B. Smith, and B. Williams. "Design of the Remote Agent Experiment for Spacecraft Autonomy." In Proceedings of the IEEE Aerospace Conference, 1998.

R. Morris, X. Descombes, J. Zerubia, M. Berthod. "Estimation of Markov Random Field prior parameters using Markov chain Monte Carlo Maximum Likelihood." To Appear in IEEE Transactions on Image Processing

R. Morris "Auxilliary Variables for Markov Random Fields with Higher Order Interactions." Submitted to Statistics and Computing.

L. Oliker, R. Biswas, S. K. Das, and D. J. Harvey, "Portable Parallel Programming for the Dynamic Load Balancing of Unstructured Grid Applications." IEEE Transactions on Software Engineering 1999, submitted. Also International Parallel Processing Symposium (IPPS99), submitted.

L. Oliker, and R. Biswas, "Dynamic Load Balancing for Parallel and Adaptive Unstructured Grid Computations." SIAM Parallel Processing Meeting, submitted.

L. Oliker and R. Biswas, "A Parallel Algorithm for Tetrahedral Mesh Adaption," Algorithmica, submitted.

L. Oliker, K. Schloegel, G. Karypis, V. Kumar, and R. Biswas, "A Performance Study of Diffusive vs. Remapped Load-Balancing Schemes." 11th International Conference on Parallel and Distributed Computing Systems, 1998, to appear.

L. Oliker, and R. Biswas, "Experiments with Repartitioning and Load Balancing Adaptive Meshes." Grid Generation and Adaptive Algorithms, (M. Luskin, J.E. Flaherty, and M. Bern, eds.), IMA Volumes in Mathematics and its Applications, Springer-Verlag, to appear.

L. Oliker, R. Biswas, and H.N. Gabow, "Performance Analysis and Portability of the PLUM Load Balancing System," Euro-Par'98 Parallel Processing, Lecture Notes in Computer Science, Springer-Verlag.

L. Oliker and R. Biswas, "PLUM: Parallel Load Balancing for Adaptive Unstructured Meshes." Journal of Parallel and Distributed Computing, August 1998.

L. Oliker, and R. Biswas, "Load Balancing Sequences of Unstructured Adaptive Grids." 4th International Conference on High Performance Computing, 1997, pp. 212-217.

L. Oliker, "PLUM: Parallel Load Balancing for Adaptive Unstructured Meshes." Ph.D. Dissertation, University of Colorado, Dept. of Computer Science, Nov. 1997.

- L. Olikar, R.C. Strawn, and R. Biswas, "New Computational Methods for the Prediction and Analysis of Helicopter Noise." *Journal of Aircraft*, Vol. 34, No. 5, 1997, pp. 665-672.
- B. Pell, P.P. Nayak, D.E. Bernard. (JPL), S. Chien (JPL), E. Gat (JPL), N. Muscettola (Recom) M.D. Wagner (Fourth Planet), and B.C. Williams (NASA Ames). "An autonomous spacecraft agent prototype." *Autonomous Robots* 5(1), 1998.
- B. Pell, and E. Gat (JPL) "Smart Executives for Autonomous Spacecraft." *IEEE Intelligent Systems*, Sept/Oct, 1998. Special Issue on Spacecraft Autonomy.
- B. Pell, G. A. Dorais (Caelum), C. Plaunt (Caelum), and R. Washington (Caelum). "The Remote Agent Executive: Capabilities to Support Integrated Robotic Agents." Invited submission to *Autonomous Robots Journal*, special issue on Integrating Robotic Architectures.
- B. Pell, P.P. Nayak, E. Gamble (JPL), E. Gat (JPL), R. Keesing (Caelum), J. Kurien (Caelum), B. Millar (Caelum), C. Plaunt (Caelum), and B.C. Williams (NASA Ames). "A hybrid procedural/deductive executive for autonomous spacecraft." *Autonomous Agents and Multi-Agent Systems* 2(1), June, 1998 (To appear).
- B. Pell, N. Muscettola (Recom), P. Morris (Caelum), and B. Smith (JPL), "Issues in temporal reasoning for autonomous control systems." Invited submission to *Autonomous Robots Journal*.
- A. Jonsson, and Smith, D., "STRIPS Planning and Temporal CSPs". In "Working Notes of Planning as Combinatorial Search Workshop", 1998.
- Susie Go, "Multilevel Methods on Unstructured Grids", Ph.D. Dissertation, Department of Mathematics, University of California Los Angeles, March 1998.
- P. C. Cheeseman, D. G. Balwin, and W. J. Emery, "Higher Resolution Earth Surface Features from Repeat Moderate Resolution Satellite Imagery." *IEEE Trans. on Geo. And Rem. Sens.*, Vol. 36, No. 1, Jan. 1998.
- R.E. Vinkhuyzen, and J.W. Whalen, "Workplace Studies: Recovering Work Practices and Informing Systems Design." Forthcoming: "Expert Systems in (Inter)Action: Diagnosing Document Machines Problems Over the Telephone." In: Chr. Heath, J. Hindmarsh, P. Luff (Eds.) Cambridge: Cambridge University Press
- S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomistic Simulation of SiH Interactions with Silicon Surfaces During Deposition from Silane Containing Plasmas", *Applied Physics Letters*, vol. 72, pp. 578-580 (1998).
- S. Ramalingam, D. Maroudas, and E. S. Aydil, "Interaction of SiH Radicals from a Silane Plasma with Silicon Surfaces : An Atomic-Scale Simulation Study", to appear in *Journal of Applied Physics*, vol. 84, pp. 3895-3911 (1998).
- S. Ramalingam, A. Lopez, D. Maroudas, and E. S. Aydil, "Computational Atomic-Scale Study of the Reactivity of Molecular Fragments from SiH₄/H₂ Plasma with Amorphous Silicon Surfaces", in *Proceedings of the International Symposium on Thin Film Materials, Processes, Reliability, and Applications: Thin Film Processes*, edited by G. S. Mathad, M. Meyyappan,

and M. Engelhardt, *The Electrochemical Society Proceedings Series*, vol. **97-30**, The Electrochemical Society, Pennington NJ (1998), pp. 35-46.

S. Ramalingam, D. Maroudas, E. S. Aydil, and S. P. Walch, "Abstraction of Hydrogen by SiH_3 from Hydrogen-Terminated $\text{Si}(001)-(2 \times 1)$ Surfaces", *Surface Science Letters*, in press (1998).

S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomic-Scale Analysis of the Reactivity of Radicals from Silane/Hydrogen Plasmas with Silicon Surfaces", in *Thin-Film Structures for Photovoltaics*, edited by E. D. Jones, R. Noufi, B. L. Sopori, and J. Kalejs, *Materials Research Society Symposia Proceedings*, vol. **485**, in press (1998).

S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomic-Scale Analysis of Plasma-Enhanced Chemical Vapor Deposition from Silane/Hydrogen Plasmas on Silicon Substrates", in *Amorphous and Microcrystalline Silicon Technology*, edited by S. Wagner, M. Hack, H. M. Branz, R. Schropp, and I. Shimizu, *Materials Research Society Symposia Proceedings*, vol. **507**, in press (1998).

S. Ramalingam, D. Maroudas, and E. S. Aydil, "Visualizing Radical-Surface Interactions in Plasma Deposition Processes: Reactivity of SiH_3 Radicals with Silicon Surfaces", submitted to *IEEE Transactions of Plasma Science* (1998).

S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomistic Simulation Study of the Interactions of SiH_3 Radicals with Silicon Surfaces", manuscript in preparation to be submitted to *Journal of Applied Physics*.

S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomic-Scale Modeling of Plasma Surface Interactions in the PECVD of Silicon", manuscript in preparation to be submitted to *The Electrochemical Society Symposia Series*.

A. Sohn, Y. Paek, J. Ku, Y. Kodama, Y. Yamaguchi, Communication Studies of Three Distributed-Memory Multiprocessors, Accepted to appear in the Proceedings of IEEE Symposium on High Performance Computer Architecture, as of October 1998. The conference will be held in January 1999. One of the high quality conferences in computer architecture.

A. Sohn and H. D. Simon, S-HARP: A Scalable Parallel Dynamic Partitioner for Adaptive Computations, in the Proceedings of ACM/IEEE Supercomputing 98, November 1998, Orlando, Florida. - The flagship conference in high performance computing.

A. Sohn and Y. Kodama, Load Balanced Parallel Radix Sort, in Proceedings of the 12th ACM International Conference on Supercomputing, Melbourne, Australia, July 1998, pp.305-312. One of the high quality conferences in high performance computing.

A. Sohn, R. Biswas, Special Issue on Dynamic Load Balancing, *Journal of Parallel and Distributed Computing* 47, December 1997, pp.99-101. One of the two flagship journals in high performance computing.

A.Sohn, H. D. Simon, and R. Biswas, HARP: A Dynamic Spectral Partitioner, Journal of Parallel and Distributed Computing 50, April 1998, pp.88-103. One of the two flagship journals in high performance computing.

V. PAPERS SUBMITTED TO REFERRED JOURNALS

BARNEY PELL

An autonomous spacecraft agent prototype, B. Pell, D.E. Bernard. (JPL), S. Chien (JPL), E. Gat (JPL), N. Muscettola (Recom), P.P. Nayak, M.D. Wagner (Fourth Planet), and B.C. Williams (NASA Ames). *Autonomous Robots* 5(1), 1998.

Remote agent: To boldly go where no AI system has gone before, N. Muscettola (Recom), P.P. Nayak, B. Pell, and B.C. Williams (NASA Ames). *Artificial Intelligence*, 103(1/2), August 1998.

Smart Executives for Autonomous Spacecraft, E. Gat (JPL) and B. Pell. *IEEE Intelligent Systems*, Sept/Oct, 1998. Special Issue on Spacecraft Autonomy. To Appear.

The Remote Agent Executive: Capabilities to Support Integrated Robotic Agents, Barney Pell, Gregory A. Dorais (Caelum), Christian Plaunt (Caelum), and Richard Washington (Caelum). Invited submission to *Autonomous Robots Journal*, special issue on Integrating Robotic Architectures.

A hybrid procedural/deductive executive for autonomous spacecraft, B. Pell, E. Gamble (JPL), E. Gat (JPL), R. Keesing (Caelum), J. Kurien (Caelum), B. Millar (Caelum), P.P. Nayak, C. Plaunt (Caelum), and B.C. Williams (NASA Ames). *Autonomous Agents and Multi-Agent Systems* 2(1), June, 1998 (To appear).

Issues in temporal reasoning for autonomous control systems, N. Muscettola (Recom), P. Morris (Caelum), B. Pell, and B. Smith (JPL). In Wooldridge, M., ed., *Proceedings of the Second International Conference on Autonomous Agents*. ACM Press, 1998.

Mission operations with an autonomous agent, B. Pell, S. Sawyer (Lockheed), N. Muscettola (Recom), B. Smith (JPL), and D.E. Bernard (JPL). In *Proceedings of the IEEE Aerospace Conference*, 1998.

Infusion of autonomy technology into space missions - DS1 lessons learned, A.S. Aljabri (JPL), D.E. Bernard (JPL), D.L. Dvorak (JPL), G.K. Man (JPL), B. Pell, and T.W. Starbird (JPL). In *Proceedings of the IEEE Aerospace Conference*, 1998.

Abstract resource management in an unconstrained plan execution system, E. Gat (JPL) and B. Pell. In *Proceedings of the IEEE Aerospace Conference*, 1998.

Design of the remote agent experiment for spacecraft autonomy, D.E. Bernard (JPL), G.A. Dorais (Caelum), C. Fry (Caelum), E.B. Gamble, Jr. (JPL), B. Kanefsky (Caelum), J. Kurien (Caelum), W. Millar (Caelum), N. Muscettola (Recom), P.P. Nayak, B. Pell, K. Rajan (Caelum), N. Rouquette (JPL), B. Smith (JPL), and B.C. Williams (NASA Ames). In *Proceedings of the IEEE Aerospace Conference*, 1998.

B. Pell, G.A. Dorais (Caelum), C. Plaunt (Caelum), and R. Washington (Caelum), The remote agent executive: Capabilities to support integrated robotic agents. In Schultz, A., and Kortenkamp, D., eds., *Procs. of the AAI Spring Symp. on Integrated Robotic Architectures*. Palo Alto, CA: AAI Press, 1998.

C. Plaunt (Caelum), Kanna Rajan (Caelum), Nicola Muscettola (Recom), and B. Pell, Integrated Planning and Execution for Satellite Tele-Communications. In *Procs. of the AAI*

Fall Symp. on Integrated Planning For Autonomous Agent Architectures, Orlando, Fl, October 23-25, 1998 (To Appear).

G. A. Dorais (Caelum), R. Peter Bonasso (JSC/Traclabs), David Kortenkamp (JSC/Traclabs), B. Pell, and Debra Schreckenghost (JSC/Traclabs), Adjustable Autonomy for Human-Centered Autonomous Systems on Mars. In Procs. of the 1998 Mars Society Conference (To Appear).

VI. SEMINARS AND COLLOQUIA

ARL JONSSON

Attended: First International Workshop on Planning and Scheduling for Space Exploration and Science. Held in Oxnard, California, October 28-30, 1997.

Attended: Planning as Combinatorial Search; a Workshop held in conjunction with AIPS98. Held in Pittsburgh, Pennsylvania, June 7.

Attended: Fourth International Conference on Artificial Intelligence and Planning Systems (AIPS98). Held in Pittsburgh, Pennsylvania, June 7-10, 1998.

Attended: Fifteenth National Conference on Artificial Intelligence (AAAI98) and Tenth Conference on Innovative Applications of Artificial Intelligence (IAAI98). Held in Madison, Wisconsin, July 26-30, 1998.

Presented: "Cyclic Scheduling", a talk at weekly ISG seminar on June 23rd, 1998.

Attended: The weekly ISG seminars which included both talks by visiting speakers and presentations by group members.

BARNEY PELL

Reviewer for International Conference on Practical Applications of Autonomous Agents and Multi-Agent Systems (PAAM-98).

Reviewer for International Computer Chess Association Journal (ICCA).

Reviewer for International Conference on Computer Games, Japan, October 1998.

Reviewer for SBIR proposal on satellite communications optimization.

AAAI Fall Symposium on Model-Directed Autonomous Systems, Boston, MA, November, 1997.

NASA Planning Workshop, Oxnard, CA, November, 1997.

Workshop on Spacecraft Autonomy, Albuquerque, NM, December, 1997.

IEEE Aerospace Conference, Snowmass, CO, March, 1998.

Second International Conference on Autonomous Agents (Agents98), Minneapolis, MI, May, 1998. Presented two papers, and gave tutorial on "Agent Architecture for Autonomous Control Systems". Approximately 50 people attended the tutorial.

Workshop on Integrated Vehicle Health Management (IVHM) for Reusable Launch Vehicles (RLV), KSC 4/21-23/98. The workshop focused on the current and future technology validation flights for IVHM. Current flights are focusing on new fiber optic sensor technologies and the distributed communication and data processing technologies necessary to manipulate the large data volumes resulting from massive sensor installations. We gave a presentation on the Remote Agent, which was received with enthusiasm, and we discussed

ways such an agent could leverage the current capabilities as we head toward more autonomous data processing.

International Robocup Workshop, Paris, France, 7/2-3/98. We gave an Invited Talk on "Future NASA Missions: Autonomous Agents, Teams, and Robocup".

International Workshop on Agent Theories, Architectures, and Languages, Paris, France, 7/4-6/98. We spoke on an Invited Panel on the subject of "The Future of the Belief-Desires-Intentions (BDI) Model for Agent Architectures".

International Conference on Multi-Agent Systems (ICMAS), Paris, France, 6/30-7/8/98.

DAVID L. GEHRT

Attended the SANS 98 conference in Monterey, CA during May, 1998.

ERIK VINKHUYZEN

Attended the American Sociological Association's annual conference in San Francisco from 8/21/98 - 8/25/98.

FRANCES JAMES

To gain a broader background in computational linguistics, project members attended the COLING-ACL '98 conference held in Montreal in August. Also, the AVIOS '98 conference in San Jose was very helpful for finding out more about commercial speech recognition software and hardware.

LEONID OLIKER

Attended 1998 Performance-Engineered Information Systems Workshop, NASA Ames Research Center September 1998.

Presented "Performance Analysis and Portability of the PLUM Load Balancing System," at Euro-Par'98, South Hampton England, September 1998.

Presented "PLUM: Parallel Load Balancing for Adaptive Unstructured Meshes," at the National Energy Research Scientific Computing Center (NERSC) in Lawrence Berkeley National Laboratory, Berkeley CA, September 1998.

Presented "PLUM: Parallel Load Balancing for Adaptive Unstructured Meshes," at the Lawrence Livermore National Laboratory, Livermore CA, September 1998.

Presented "Experiments with Repartitioning and Load Balancing Adaptive Meshes," at the Institute for Computer Applications in Science and Engineering (ICASE) at NASA Langley Research Center, Langley VA, August 1998.

Attended the Symposium on Solving Irregularly Structured Problems in Parallel (IRREGULAR 98), Berkeley CA, August 1998.

Presented "PLUM: Parallel Load Balancing for Adaptive Unstructured Meshes," Thesis Defense, University of Colorado, Boulder CO, November 1997.

PANDURANG NAYAK

During the year we presented Colloquia on the Remote Agent and the Livingstone Model-based diagnosis and recovery system at Stanford University, University of California Berkeley, USC/ISI, and Sonoma State.

PETER CHEESMAN

Keynote speaker at the 9th. Australian A.I. Conference, Brisbane, Australia, July 1998.

Invited speaker, MaxEnt Conference, Garching, Germany, July 1998.

P. Nayak, B. Pell, N. Muscettola, and B. Williams. "Remote Agent: To Boldly Go Where No AI System Has Gone Before." To appear in the Artificial Intelligence Journal, 103(1/2), August 1999.

Keynote speaker at the 9th. Australian A.I. Conference, Brisbane, Australia, July 1998.

Invited speaker, MaxEnt Conference, Garching, Germany, July 1998.

ROBIN MORRIS

Attended the SPIE conference "Bayesian Inference for Inverse Problems" San Diego, July 1998

Gave an invited seminar at Trinity College, Dublin, Ireland, November 1997 and at INRIA, Sophia Antipolis, France, September 1998

S. RAMALINGAM

S. Ramalingam, D. Maroudas, and E. S. Aydil, "Computational Atomic-Scale Study of the Reactivity of Molecular Fragments from SiH_4/H_2 Plasma with Amorphous Silicon Surfaces", The 1997 Joint International Meeting of the Electrochemical Society and the International Society of Electrochemistry, September 1997, Paris, France.

S. Ramalingam, D. Maroudas, and E. S. Aydil, "Interactions of Molecular Fragments from Silane/Hydrogen Plasma with Silicon Surfaces: An Atomic-Scale Computational Study", 44th Annual Symposium of the AVS, October 1997, San Jose, California.

S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomic-Scale Analysis of the Reactivity of Molecular Fragments from Silane/Hydrogen Plasmas with Surfaces of Hydrogenated Amorphous Silicon Films", AIChE Annual Meeting, November 1997, Los Angeles, California.

D. Maroudas, S. Ramalingam, and E. S. Aydil, "Atomic-Scale Modeling of Interactions Between Plasmas and Silicon Surfaces and Plasma-Enhanced Chemical Vapor Deposition of Silicon Films", AIChE Annual Meeting, November 1997, Los Angeles, California.

S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomic-Scale Analysis of the Reactivity of Radicals from Silane/Hydrogen Plasmas with Silicon Surfaces", MRS Fall Meeting, December 1997, Boston, Massachusetts.

D. Maroudas, "Structure, Surface Reactivity, and Reliability of Metallic and Semiconductor Thin Films: A Computational Materials Science Study", Silicon Science and Technology Division Colloquium, IBM T. J. Watson Research Center, December 1997, Yorktown Heights, New York.

D. Maroudas, "Structure, Surface Reactivity, and Reliability of Metallic and Semiconductor Thin Films: A Computational Materials Science Study", Computational Materials Physics Seminar, Department of Physics and Division of Engineering and Applied Sciences, Harvard University, December 1997, Cambridge, Massachusetts.

D. Maroudas, "Structure, Surface Reactivity, and Reliability of Metallic and Semiconductor Thin Films: A Computational Materials Science Study", Chemical Engineering Departmental Seminar, Princeton University, February 1998, Princeton, New Jersey.

S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomic-Scale Analysis of Surface Reaction Mechanisms in the Plasma-Enhanced Chemical Vapor Deposition of Silicon Thin Films", APS March Meeting, March 1998, Los Angeles, California.

D. Maroudas, "Structure, Surface Reactivity, and Reliability of Metallic and Semiconductor Thin Films: A Computational Materials Science Study", Chemical Engineering Departmental Seminar, University of Minnesota, April 1998, Minneapolis, Minnesota.

D. Maroudas, "Structure, Surface Reactivity, and Reliability of Metallic and Semiconductor Thin Films: A Computational Materials Science Study", Chemical Engineering Departmental Seminar, Massachusetts Institute of Technology, April 1998, Cambridge, Massachusetts.

S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomic-Scale Analysis of Plasma-Enhanced Chemical Vapor Deposition from Silane/Hydrogen Plasmas on Silicon Substrates", MRS Spring Meeting, April 1998, San Francisco, California.

D. Maroudas, "Structure, Surface Reactivity, and Reliability of Metallic and Semiconductor Thin Films: A Computational Materials Science Study", Chemical Engineering Departmental Colloquium, University of Illinois at Urbana-Champaign, September 1998, Urbana-Champaign, Illinois.

D. Maroudas, "Structure, Surface Reactivity, and Reliability of Metallic and Semiconductor Thin Films: A Computational Materials Science Study", Mechanical & Environmental Engineering Departmental Seminar, University of California, Santa Barbara, October 1998, Santa Barbara, California.

D. Maroudas, S. Ramalingam, and E. S. Aydil, "Atomic-Scale Modeling of Plasma-Surface Interactions in the PECVD of Silicon", Symposium on Fundamental Gas-Phase and Surface Chemistry of Vapor-Phase Synthesis, The 194th Meeting of the Electrochemical Society, November 1998, Boston, Massachusetts.

D. C. Marra, S. Ramalingam, E. Edelberg, D. Maroudas, and E. S. Aydil, "Surface Reactions and Hydrogen Coverage on Plasma Deposited Hydrogenated Amorphous Silicon and Nanocrystalline Silicon Surfaces", 45th Annual Symposium of the AVS, November 1998, Baltimore, Maryland.

S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomistic Simulation of Plasma Enhanced Chemical Vapor Deposition of Hydrogenated Amorphous Silicon Films", 45th Annual Symposium of the AVS, November 1998, Baltimore, Maryland.

D. Maroudas, "Structure, Surface Reactivity, and Reliability of Metallic and Semiconductor Thin Films: A Computational Materials Science Study", Chemical Engineering Departmental Colloquium, Stanford University, November 1998, Stanford, California.

S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomic-Scale Analysis of Surface Reaction Mechanisms during Plasma Deposition from Silane Discharges on Silicon Substrates", AIChE Annual Meeting, November 1998, Miami Beach, Florida.

S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomic-Scale Analysis of Plasma-Enhanced Chemical Vapor Deposition of Si Films from Silane/Hydrogen Discharges", AIChE Annual Meeting, November 1998, Miami Beach, Florida.

A. Sohn, "Supercomputing (ICS'98), Melbourne, Australia, July 9, 1998." Research paper at the ACM International Conference.

A. Sohn, presented a research paper at the Irregular'98, August 10, 1998, Berkeley, California.

A. Sohn, attended the ACM/IEEE Conference on Parallel Architectures and Compilation Techniques (PACT'98), October, 1998, as one of the conference organizers (publicity chair), a program committee member, and the conference general chair-elect for PACT'2000.

A. Sohn, other publications, visits and presentations that are related to the IPG project and high performance computing but these were by NSF, NJIT, and other funding agencies.

Abbreviations AIChE: American Institute of Chemical Engineers
APS: American Physical Society
AVS: American Vacuum Society
ECS: Electrochemical Society
MRS: Materials Research Society

VII. OTHER ACTIVITIES

PANDURANG NAYAK

Associate Editor of the Journal of Artificial Intelligence Research.

Co-chair of the Ninth International Workshop on Principles of Diagnosis.

Member of the Program Committee for AAAI-98 and the Workshop on Qualitative Reasoning about Physical Systems (QR98).

Co-taught a tutorial on "Model-based Autonomous Systems" at AAAI-98.

Taught a class at Stanford University on "Reasoning Methods in Artificial Intelligence.

Appointed as a Consulting Faculty in the Department of Computer Science at Stanford University.

MARJORY J. JOHNSON

Co-organized NASA/NREN Quality of Service Workshop, August, 1998

Member of the Network Research Team for the Large Scale Networking committee

Member of program committees for IFIP Conference on High Performance Networking (HPN '98), European Conference on Multimedia Applications, Services and Techniques (ECMAST '98), and Super Computing '99

Reviewed several research proposals for NSF

Attended IEEE INFOCOM '98 Conference, March/April, 1998

Attended IEEE 1998 Sixth International Workshop on Quality of Service, May, 1998

Attended Internet2 QoS Workshop, May, 1998 and Internet2 Fall Project Meeting, September, 1998

Attended ACM SIGCOMM '98 Conference, September, 1998

Member of ISO/T C20/USSCAG 13 committee to develop communication standards for space missions

DAVID L. GEHRT

Part of the beta program for the RedHat 5.2 version of Linux on the Sun Sparc architecture.

Other than these things we continued our past efforts to assist the network operations and Center UNIX system administrators in dealing with problems that arise from time to time.

Was asked to work on two Documents for the Center. One is a review of the Center computer and network security policy, and the other is a rewrite of the Computer security incident handling procedure. This work was set aside by other taking, but has recently resurfaced as a priority.

ROBIN MORRIS

Acted as a reviewer for the International Journal of Computer Vision

LEONID OLIKER

On leave-of-absence for three months 1/15/98-4/15/98.

BARNEY PELL

Worked with our research team on several major applications of the agent research.

1. Remote Agent Experiment (RAX)
2. The development and flight preparation of the Remote Agent Experiment (RAX), an autonomy experiment to be demonstrated on Deep Space One (DS1), the first deep space technology validation flight in NASA's New Millennium Program (NMP). The RAX will demonstrate planning, scheduling, execution, diagnosis, and recovery in flight, and validate the technology for use in future space missions. In this year, achieved all scheduled major milestones on the preparation path for DS1 launch in late October. As the RAX experiment itself will not run until mid-1999, the major goals before DS1 launch were to address the system design and interactions with the DS1 flight software and ensure that the RAX-manager software (which will launch with DS1) would launch with interfaces and resources that are sufficient to support the later RAX experiment. To this end, integrated the major RAX components, demonstrated the system on nominal and failure scenarios running on a flight-like processor, designed and executed test scenarios on the subsystems and on the integrated system, conducted a test review for the integrated RAX components, and finally completed a launch readiness review.
3. Remote Agent for Upcoming Space Science Missions
4. Second, following on the success of the DS1 Remote Agent, started addressing new challenges for the Remote Agent applied to science spacecraft in upcoming missions.
5. RA for Interferometers
6. In June, initiated a tight collaboration with JPL to help engineers apply the Remote Agent to Interferometers. Interferometers are the key scientific instruments for NASA's Origins program, supporting such missions as the Space Interferometry Mission (SIM), DS-3, Planet finder, and the Keck Observatory upgrade. Three months later, in September, demonstrated RA (the EXEC and MIR components, no planning yet) performing closed-loop control of the interferometer hardware within the main interferometer testbed at JPL. This required integration between internal RA components within a new testbed, interfacing between RA and an existing socket-based customer interface, and an RA framework for continuing interferometer model and sequencing development. The demonstration was provided to JPL engineers and managers involved in interferometry missions and other missions.
7. This accomplishment is significant for several reasons. This is the first time RA has controlled actual hardware. Demonstrations like this increase confidence and have led to

RA being base lined for a variety of interferometry missions. This is also the first time RA is being pulled and partially applied by customers, and demonstrates that the RA is becoming increasingly portable. This also demonstrates new application areas for RA, to control specific complex instruments rather than an entire spacecraft.

8. X2000 and other advanced operations architectures
9. Another collaboration, also with JPL, is the X2000 Mission Data System project, which is attempting to design a new mission architecture to support multiple deep space missions. Our group participated in the technical design of the software architecture, based on lessons from the Remote Agent and our future technology directions. Also, explored collaborations with KSC, JSC, and GSFC to support increasing capability of their missions using autonomous agent technology.

10. Communication Satellite Network Management

11. A third research application area is in communication satellite networks. Investigated how autonomous capabilities could help these systems by increasing reliability and resource utilization. In this period, we initiated and carried out a project with Lockheed to develop a method for advanced resource reservations and on-board dynamic load management. Completed the initial design and implementation of a prototype, which demonstrates autonomous antenna (beam) movement to provide satellite coverage over areas requested in advance, migrating calls from one beam to another to support such movements while maintaining load, and cancellation of lower priority calls to preserve quality-of-service (QoS) guarantees. Presented a demonstration of the Call Admission Controller (CAC) we developed, as part of the Milsatcom Advanced EHF Program Modeling & Simulation Demo Day, at Lockheed-Martin, 6/18/98. According to Lockheed collaborators, this was "a highly effective demonstration of the improved efficiency in link resource allocation due to autonomous planning and real-time policy execution. This demo represented an intense effort in a demanding schedule and its timely completion added significant capability to the day's demonstrations. The Call Admission Controller fills an important analysis role in the system design phase, especially in the context of the Modeling & Simulation Lab. In addition, the technology holds promise for migration to an operational setting."

12. Mobile Robotics

13. A fourth application area is extending the Remote Agent to address the needs of mobile robots. We are currently developing an onboard resource management capability for planetary rovers. The resource manager will enable new levels of autonomy for Mars rovers, as they will be able to manage power, communications, and thermal resources, interrupting lower-priority activities to ensure that critical activities get the resources they need. We are also developing a new sequence execution system, which accepts conditional (branching and event-based) sequences and backup plans. These new capabilities will let ground operators build more challenging mission sequences, getting increased science return and robustness and also increasing the useful lifetime of planetary rovers. In the current period, met with researchers in robotics (the Intelligent Mechanisms Group at NASA Ames, and the Mars '03 Rover team at JPL) to determine requirements. Then designed languages and application programmer interfaces (API) for both the Resource Manager and the Conditional Executive, developed fundamental algorithms and then implemented prototypes of both systems to support the required capabilities. These capabilities are now being integrated into a next-generation Rover Executive, which will be demonstrated during a field test toward the end of calendar year 1998.

14. RIACS wrote a division-level proposal on "Human-centered autonomy", describing ideas for developing an agent which can minimize the necessity for human monitoring and guidance, but also maximize the opportunity for human interaction. This involves supporting multiple levels of commanding and monitoring across different activities and mission phases. Published a paper on "Mission Operations with an Autonomous Agent" which introduced many of these issues to the research community, published a paper on adjustable autonomy in future Mars Habitats, formed collaborations with researchers at Honeywell, Lockheed, BOEING, and JSC who are interested in this general area, and organized a AAAI Spring Symposium on "Agents with Adjustable Autonomy" to be held next spring.
15. Progress was made in the area of real-time guarantees for agents, and developed a new algorithm for plan execution that guarantees a maximum latency. We published this work in an international conference and journal.
16. An investigation of fundamental agent capabilities was started. RIACS commenced work on the architecture of the next-generation Remote Agent and a modular, generic Executive, starting with the resource management and plan execution capabilities. Also wrote division-level proposals on "Independent Deployable Execution Agents (IDEAS)", and "On-board autonomy for rovers", both based on the work described above. The IDEAS work is now submitted for journal publication.

BARNEY PELL

TUTORIALS:

Dr. Barney Pell presented a tutorial on "Agent Architecture for Autonomous Control Systems" at the Second International Conference on Autonomous Agents, in Minneapolis, MN, 5/09/98. The tutorial provided a survey of the work in the field on different architectures for building autonomous control systems, which integrate such agent capabilities as planning, scheduling, execution, real-time control, monitoring, and diagnosis. The presentation discussed the importance of architecture in addressing key themes including knowledge acquisition, integration, robustness, modularity, reuse, uncertainty, and verification. Particular emphasis was given to significant real-world applications, including mobile robotics, factory automation, and autonomous spacecraft.

RESEARCH VISITS/PRESENTATIONS:

Visiting Researcher, Electrotechnical Laboratory (ETL), Tsukuba, Japan, 2/23--3/8/98. During this trip we gave presentations on Autonomous Agent architecture at ETL, Japanese space agency (NASDA), Sony Computer Science Laboratory, and Hitachi Systems Development Laboratory.

Dr. Barney Pell gave an Invited Presentation at Sony's Computer Science Lab (CSL) 10th Anniversary Symposium, in Tokyo, Japan, May 28th, 1998. The topic of the talk was "Remote agent: To boldly go where no AI system has gone before". The talk, which was attended by over 400 people, was based on the Invited Talk given at IJCAI97 by N. Muscettola, P. Nayak, B. Pell, and B.C. Williams. During the visit, Pell also attended Sony CSL's Open House to learn about projects going on in their computer laboratory.

Dr. Barney Pell visited the Honeywell Technology Center (HTC), in Minneapolis, MN, on 5/14/98. Pell gave a presentation on Autonomous Systems Research conducted at NASA Ames IC Division and met with researchers working on formal validation, real-time operating

systems, planning and scheduling, and management of mission-critical control systems in abnormal situations. The visit identified a number of research thrusts which are areas for future collaboration between HTC and NASA Ames.

Dr. Barney Pell visited NASA Kennedy Space Center, 4/21-23/98. During the visit, we explored opportunities for technical collaboration in the area of model-based reasoning and autonomous systems between ARC and KSC. We met with technical managers and the researchers developing the KATE model-based reasoning system. In the short term, we considered the idea of working jointly on an In-Situ Propellant Production (ISPP) scenario, building models in both languages to compare and contrast the capabilities.

Met with visiting researchers from GSFC, 5/6/98 and also in 6/8/98, for technical interchange. We identified possible collaborations in the areas of spacecraft constellations and formation flying, in the next-generation hubbell telescope, and in the sun-earth connection. A follow-up visit at GSFC is being planned for October, 1998.

ORGANIZATION AND LEADERSHIP:

Co-Lead, Autonomy Integrated Product Development Team for NASA's New Millennium Program (NMP).

Co-chair of AAAI Spring Symposium on Agents With Adjustable Autonomy, March 1999.

Member, Organizing Committee, International Conference on Practical Applications of Autonomous Agents and Multi-Agent Systems (PAAM-98).

Member, Organizing Committee, International Conference on Practical Applications of Autonomous Agents and Multi-Agent Systems (PAAM-99).

Member, Organizing Committee, International Conference on Computer Games, Japan, October 1998.

RESEARCH AWARDS:

NASA Group Achievement Award (Team Award), NASA Ames Research Center, September 1998. For development of the Remote Agent Experiment for New Millennium Deep Space One.

ARI JONSSON

Reviewing duties performed:

For the 15th National Conference on Artificial Intelligence (AAAI98). As a member of the program committee, reviewed nine articles in the fields of constraint satisfaction, search and scheduling.

For NASA's Small Business Innovative Research (SBIR) program. Reviewed five proposals for artificial intelligence research related to both space and aeronautics.

Rewards received:

Ari Jonsson, Pandurang Nayak, Barney Pell

Three research staff members at RIACS (USRA's Research Institute for Advanced Computer Science) shared in a 1998 Ames Honor Award. Pandurang Nayak, Barney Pell and Ari Jonsson were part of the Ames Deep Space 1 Remote Agent Team whose work was recognized.

The team citation noted "The Ames Deep Space 1 Remote Agent Team has successfully produced the most ambitious software system ever created to control high level spacecraft functions...Ames technologies were selected for flight validation on the first New Millennium spacecraft, and asteroid and comet fly by mission named Deep Space 1."

Deep Space 1 was launched October 15, 1998.

VIII. RIACS STAFF

ADMINISTRATIVE STAFF

Robert C. Moore, Director – Ph.D, Artificial Intelligence, Massachusetts Institute of Technology, 1979. Using natural-language-based knowledge sources in speech recognition (4/13/98 – present)

Joseph Oliger, Associate Director - Ph.D., Computer Science, University of Uppsala, Sweden, 1973. Numerical Methods for Partial Differential Equations (03/25/91 – 9/30/98).

Diana Martinez, Administrator (8/18/97 - present)

Beatrice Burnett, Administrative Assistant (11/3/97 - present)

Steven Suhr, Systems Administrator (11/1/96 - present)

SCIENTIFIC STAFF

Esfandiar Bandari, Ph.D, 1995, Computer Science, University of British Columbia, Computational Vision, Signal Processing, 3-D Reconstruction and Medical Imaging (9/2/98 – present)

Peter J. Cheeseman, Ph.D., 1979, Monash University, Artificial Intelligence, computational complexity, bayesian inference, computer vision, plasma physics (9/1/97-present)

Rei J. Cheng, BS, 1973, Pharmacy, Taipei Medical College, 3-D scientific visualization and medical visualization (9/1/98 – present)

Dave Gehrt, JD Law, University of Washington, 1972, UNIX system administration, security, and network based tools (1/84 - 7/85, 2/1/88 - present).

Frances H. James, Ph.D., 1998, Computer Science, Stanford University, Speech understanding technologies and applications, developing interfaces using speech. (7/13/98 – present)

Marjory J. Johnson, Ph.D., 1970, Mathematics, University of Iowa, High-performance networking for both space and ground applications (1/9/84 -present).

Ari K. Jonsson, Ph.D., 1997, Computer Science, Stanford University, New Techniques in constraint satisfaction and Automated Planning. (10/1/97 - present)

Robin D. Morris, Ph.D., 1995, Engineering, Trinity College, Cambridge, U.K., Bayesian Inference and Image Analysis, Mathematical Modeling of Natural Imager and Associated Computational Methods. (3/1/98 – present)

Pandurang Nayak, Ph.D., 1992, Computer Science, Stanford University, 1992, Autonomous systems, model-based reasoning, diagnosis and recovery, qualitative and causal reasoning, abstraction and approximation in knowledge representation and reasoning. (11/15/97 – present)

Barney Pell, Ph.D., 1993, Computer Science, Cambridge University, Autonomous Agent Operations of control robotic and software control systems. (11/1/97 - present)

Maarten Sierhuis, MS, 1986, Engineering, Mague Polytechnic University, Intelligent multi-agent simulation agent-oriented programming languages (4/1/98 – present)

Ian A. Twombly, Ph.D., 1997, Biophysics, John Hopkins University, 1997, ????? (9/1/98 – present)

R. Erik Vinkhuyzen, Ph.D, 1997, Psychology, Zurich University, study human aspects of technology use (4/6/98 – present)

POST-DOCTORAL SCIENTISTS

James Reuther, Ph.D. - University of California, Davis, numerical optimization aerodynamic shape optimization numerical analysis CFD(4/30/96 – 9/30/98).

Leonid Oliker, Ph.D. - University of Colorado, compilation of data parallel programs (9/1/94 - present).

Susie Go, Ph.D. – 1998, Applied Math, University of California, Los Angeles, multilevel methods on unstructured grids (7/1/98 - present).

RESEARCH ASSOCIATES

Susie Go - Applied Math, University of California, Los Angeles, multilevel methods on unstructured grids (1/21/96 – 6/30/98).

STUDENTS

Shyam Ramalingam - Chemical Engineering, University of Santa Barbara, Computer Modeling of desposition of Si Films using PECVD. (9/2/97-present)

Roy Lue – Internet Based TCAD using IPG, Stanford Unversity, Stanford, California, Center for Integrated Systems. (1/1/98 – 9/30/98)

VISITING SCIENTISTS AND CONSULTANTS

Andrew Sohn, Ph.D. - Assistant Professor, New Jersey Institute of Technology, Parallel computing, high performance computing, distributed-memory with multi-processor architectures, dynamic load balancing for adaptive computations fast mesh partitioning (6/13/98-1/15/99).

David Zingg, Ph.D. - Associate Professor, University of Toronto, Canada, Development and analysis of high-accuracy numerical methods applicable to simulations of fluid flows, acoustic waves and electromagnetic waves (6/21/99-8/27/99).

Dimitrios Maroudas, - Asst. Professor, Chemical Engineering, U.C. Santa Barbara, Theoretical and Computational Materials Science with emphasis on surface science and microstructure evolution in semiconductors, metallic thin films and structural alloys (9/22/97-9/11/98)

Eli-Turkel, Ph.D. - Professor, Tel Aviv University, Algorithms for Navier-Stokes equations espiciall preconditioning for low speed flow high order accurate schemes with applications acoustics and CEM (2/4/98 – 2/17/98)

Gabor Toth, Ph.D. – Post Doctoral Scientist, Eotvos University, Development of new algorithms for magnetohydrodynamics that can be applied to astrophysical problems. (8/3/98 – 8/15/98)

Henry Tirri, Ph.D. – Professor, University of Helsinki, Finland, Probabilistic (Bayesian) and information-theoretic modeling of complex systems; stochastic optimization for complex optimization problems. (6/19/98 – 8/10/98)

Marsha Berger, Ph.D. - New York University, Computational fluid dynamics; parallel computing (1/1/93 - present).

Richard G. Johnson, Ph.D. - Physics, Indiana University, 1956, Global environmental problems and issues (11/1/92 - present).

Neil Sandam, Ph.D. - Lecturer, Queen Mary & Westfield College, England, Direct numerical simulation of transitional and turbulent fluid flow, turbulence modeling (12/30/97 – 1/6/98).

James Sethian, Ph.D. - Professor, University of California, Berkeley, Computational fluid mechanics, image processing, robotics and material sciences (3/1/97 - present)

Remi Abgrall, Ph.D. – Professor, University of Bordeaux, France, Numerical implementation of boundary conditions for first order Hamilton Jacobi equations. (7/27/98 – 8/28/98)

Ronald Henderson, Ph.D. - Sr. Scientist, California Institute of Technology, Computational Fluid Dynamics, parallel computing, hydrodynamic stability, turbulence (2/17/98-3/6/98)

Steven Senger, Ph.D. – Professor, University of Wisconsin, LaCrosse, Filtering and visualization of CT and Ultrasound data sets (9/1/98 to present)

Sankaran Venkateswaran, Ph.D. -Sr. Research Associate, University of Tennessee Space Institute, preconditioning methods and use of neural networks for application to semi-conductor materials (5/18/98 – 8/7/98)